Copy-Protection with Correlated Challenges: Point Functions and More via Decisional Coset Monogamy

Amit Behera* Alper Çakan† Vipul Goyal‡

Abstract

Copy-protection is one of the main applications of quantum information in cryptography. In copy-protection, we encode a functionality in a reusable quantum state so that it cannot be split into two states (called *freeloader adversaries*) that remain simultaneously useful. Despite a long line of research, previous works have only been able to show security with respect to *independently sampled challenges* in the plain-model. However, arguably a more natural security notion considers the two freeloader adversaries receiving the same challenge. This so-called *identical-challenge* security notion is also connected to other fundamental quantum cryptographic primitives such as unclonable bits (i.e. unclonable encryption) and copy-protection of point functions.

In this work, first we make progress on the definitional foundations of these primitives, and then prove security in the plain model for our new stronger definitions, in particular also resolving the question of copy-protection with identical challenges and copy-protection of point functions. In more detail, we obtain the following results.

- **Copy-protecting decryption keys (Single-decryptor encryption).** We define a new natural security notion for single-decryptor encryption (SDE) called *correlated challenge security*, and show that implies all previous security definitions for SDE, including identical-challenge security. Then, we prove that, assuming indistinguishability obfuscation (iO) and one-way functions, the SDE construction of Kitagawa and Yamakawa (TCC'25) satisfies correlated challenge security. We also provide an almost complete characterization of the relationship among previous SDE security notions.
- **Copy-Protecting General Functionalities with Correlated Challenges.** We define correlated challenge unclonable puncturable obfuscation (UPO), allowing arbitrary correlations among challenge points and puncturing bits, plus auxiliary information before and after splitting. Security requires only conditionally uniform bits and λ^c average conditional min-entropy in each point separately, for any constant $c > 0$; thus, in particular the challenge points may be identical. Assuming polynomially secure post-quantum iO and quantum-hard LWE, we construct correlated UPO for arbitrary polynomial-size keyed circuits with input length at least λ^c, answering the open question of Ananth, Behera, Huang, Kitagawa, Yamakawa (EUROCRYPT'26) and Çakan-Goyal (EUROCRYPT'26).
- **Applications.** Our results yield the first plain-model copy protection for point functions, k-point functions, and compute-and-compare programs under natural security definitions, and identical-challenge copy protection for general puncturable functionalities.

The technical core of our results is a new *decisional monogamy theorem for coset states*, which both simplifies the proofs and generalizes the results of existing copy-protection constructions, which may be of independent interest.

Keywords: Quantum cryptography, Copy-protection, Unclonable encryption.

*NTT Research amitbehera1767@gmail.com

†Carnegie Mellon University. alpercakan98@gmail.com.

‡NTT Research & Carnegie Mellon University. vipul@vipulgoyal.org

Contents

1 Introduction **4**
- 1.1 Our Results . . . 6
- 1.2 Organization . . . 9

2 Technical Overview **10**
- 2.1 Our Solution: Identical-Challenge Decisional Coset Monogamy of Entanglement . . . 11
  - 2.1.1 Information-Theoretic Security Proof . . . 12
  - 2.1.2 Computational Decisional Hidden-Coset Monogamy . . . 14
- 2.2 Correlated-Challenge Single-Decryptor Encryption . . . 14
- 2.3 Copy Protection from Correlated UPO . . . 16

3 Preliminaries **18**
- 3.1 Notation and Conventions . . . 18
- 3.2 Pseudorandom Functions . . . 19
- 3.3 Digital Signature Schemes . . . 20
- 3.4 Indistinguishability Obfuscation . . . 21
- 3.5 Subspace-Hiding Obfuscation . . . 22
- 3.6 Universal Hash Functions and Average Conditional Min-Entropy . . . 22
- 3.7 Injective/Lossy Functions . . . 23
- 3.8 Unclonable Encryption . . . 23
- 3.9 Quantum Information . . . 24

I Definitional Work **27**

4 Single-Decryptor Encryption Security Definitions **27**
- 4.1 Prior Definitions . . . 27
- 4.2 Correlated-Challenge Decision Security . . . 30
- 4.3 Hierarchy of Definitions . . . 32
  - 4.3.1 Summary of the Main Relations . . . 45

5 Unclonable Puncturable Obfuscation Security Definitions and Copy-Protection of General Functionalities **46**
- 5.1 Unclonable Puncturable Obfuscation definitions . . . 46
  - 5.1.1 Prior Definitions:Generalized UPO Security . . . 46
  - 5.1.2 Correlated-Challenge Security . . . 47
- 5.2 Copy-Protection Security Definitions . . . 50
  - 5.2.1 Oracular security notions . . . 54

II Quantum Information Tools **56**

6 A New Decisional Coset Monogamy-of-Entanglement Game **56**
- 6.1 Unclonable indistinguishable security of the general class of encryption schemes (Proof of Lemma 9) . . . 61
  - 6.1.1 Proof of Lemma 11 . . . 64

7 Computational Monogamy-of-Entanglement Games **67**

III Copy-Protection Constructions **73**

8 Single-Decryptor Encryption Construction **73**
8.1 Construction 73
8.2 Correctness 74
8.3 Proof of security 74

9 Unclonable Puncturable Obfuscation Construction **77**
9.1 Construction 77
9.2 Proof of Security 78
9.2.1 Formal proof of the moreover part of Theorem 23 84

10 Applications of Correlated UPO to Copy-Protection **87**

1 Introduction

Starting with the seminal work of Wiesner [Wie83], a long line of research has shown that quantum information can help us achieve remarkable cryptographic guarantees that are simply impossible with classical information. A central example is the notions of quantum money ([Wie83, AC12]) and *quantum copy protection* [Aar09]. In the latter, a functionality is encoded into a quantum state that can be evaluated, but no efficient adversary can split the state into two programs/adversaries (also called *freeloaders*) that both remain useful. The functionality can be a cryptographic object (such as a decryption key for an encryption scheme) or can be a piece of software (such as a video game or a spreadsheet program). Note that a classical program or key can always be copied bit for bit, so copy protection is indeed an inherently quantum security notion.

There has been significant recent progress on constructing copy protection for general classes of functionalities and laying down the definitional foundations of copy protection [CMP20, ALL+21, CLLZ21, AKL+22, AKL23, AB24, ABH+26, ÇG26, KY25]. Nevertheless, as we discuss in more detail below, the state of affairs remains quite unsatisfactory on both the construction and definitional fronts. Even for some of the simplest functionalities, such as *point functions*, we still do not have a provably secure copy protection scheme (under a natural security definition) in the plain model. On the definitional side, we do not have good security definitions for even copy-protecting decryption keys, which is one of the most fundamental applications of copy protection. Further, some of the existing natural security definitions for decryption keys still have not been achieved. In this work, we consider these issues and make significant progress on both construction and definitional fronts. We now discuss the issues in more detail.

Copy-protecting decryption keys. The earliest and perhaps the most important primitive in cryptography is *encryption*, and similarly, one of the earliest and most fundamental questions in quantum copy protection has been the question of copy-protecting decryption keys. Thus, this question has been one of the most studied questions in copy protection [GZ20, CLLZ21, AK21, AKL+22, KN23, AKL23, KY25] and even has a dedicated name: *single-decryptor encryption* (SDE) [CLLZ21].

In SDE, we encode a decryption key in a quantum state, and we require that no adversary can split it into two *useful* keys. Note that the encryption key and the ciphertexts are still classical. In the case of classical encryption, the notion of a *useful* adversary has been settled decades ago: The seminal work of Goldwasser and Micali [GM82, GM84] introduced the notion of indistinguishability security (now also called *CPA security*), where the adversary receives an encryption of either m_0 or m_1 with probability $1/2$, and security is considered violated when it can guess which case it is with probability non-negligibly better than $1/2$. This security notion is the *gold standard* for encryption, and also implies various other properties that one would expect from an encryption, such as unpredictability security (also called search security), where an adversary tries to guess the message from its ciphertext (the message comes from an unpredictable distribution).

However, it turns out that the question of defining security for single-decryptor encryption (SDE) is much trickier: There are 7 different security notions currently in the literature! To make things worse, even the relationships among these security notions are unclear, and there is no *gold standard* definition: No single definition is known to imply all others. For example, unlike classical encryption, it is not even known if the existing indistinguishability security implies unpredictability security (in fact, in this work we show that in general, it does not). Further, while some of the seven definitions have been achieved; for some natural ones we still do not even have provably secure classical-ciphertext constructions under standard assumptions, such as *identical-challenge* security where the two freeloaders receive the same ciphertext. This directly models a single ciphertext

distributed to many users—such as an encrypted television broadcast or a fixed encrypted file sent to all subscribers—on which two copies of the decryption key should not both remain useful. Introduced by Georgiou and Zhandry [GZ20], achieving identical-challenge security notion has proved difficult and remained open despite many attempts [CLLZ21, AKL23, CGLZR23, KN23, CHV23, AKY25a, KY25, AB24, ABH+26].

Thus, in this work, we ask the following questions.

> *Can we define a* gold-standard *security notion for single-decryptor encryption that captures the natural security intuition and uses cases, and implies existing definitions (including identical-challenge security)?*
>
> *Can we achieve a construction with provable security for all known definitions including identical-challenge security? In fact, can we achieve a construction that provably satisfies the gold standard definition, and hence as a result, all known definitions?*

Copy-protecting general functionalities. Moving beyond encryption, various works have considered copy-protecting general classes of functionalities [CMP20, ALL+21, CLLZ21, AKL+22, AKL23, AB24, ABH+26, ÇG26]. The state of the art consists of the recent concurrent works of Çakan and Goyal [ÇG26] and Ananth, Behera, Huang, Kitagawa, and Yamakawa [ÇG26, ABH+26], which show how to copy-protect any *puncturable functionality* in the plain model.

Unfortunately, these results also suffer from the same issues as the case of single-decryptor encryption. First, there is still not a single security definition for copy-protecting general functionalities. Further, these works only achieve what is called *independent-challenge security*. However, a competing natural definition called *identical-challenge security* has not been achieved and was left as an open question by both works [ÇG26, ABH+26]. This latter notion, identical-challenge security, has many practical security motivations, and aside from that, it is also tightly related to other fundamental questions, such as *unclonable bits* or *unclonable encryption* [BL20], and *copy-protecting point functions* [Aar09]. In contrast, independent-challenge security does not imply any of these notions. More broadly, prior works have struggled to tackle the question of correlation in challenge distributions and to get feasibility for correlated challenge distributions, which raises the following question.

> *Can we achieve copy protection with identical-challenge security, and more generally correlated challenge security for general classes of functionalities?*

Point functions are perhaps the simplest nontrivial programs and have natural applications such as digital lockboxes and password-verification software. Their conceptual simplicity and concrete applications led Aaronson to single out their copy protection when introducing quantum copy protection in 2009 [Aar09]; he proposed two candidate schemes in the structured ideal oracle model but left proving security under standard assumptions open. Although Coladangelo, Majenz, and Poremba later achieved security in the quantum random-oracle model [CMP20], the input challenge distribution they considered for copy protection was not the natural challenge distribution.[1] Hence, getting copy protection for point functions with the natural challenge distribution in the plain model still remains open.

Thus, in this work, we ask the following questions.

[1] Here, by natural challenge distribution, we mean that for a point function P_y with probability $1/2$, give the point y to the freeloaders $\mathcal{B}$ and $\mathcal{C}$, and with the rest of the probability, give a random point x to both the freeloaders.

Can we achieve copy protection for fundamental functionalities such as point functions and compute-and-compare functions?

In [AB24], the authors provided a modular pathway to achieving copy protection for general functionalities, via an unclonable variant of obfuscation called *Unclonable Puncturable Obfuscation (UPO)*. However, the UPO security notion itself is sensitive to challenge distributions, and [AB24] considered two variants of UPO security, one with identical challenge distribution and the other with independent challenge distributions, but each with its own set of applications to copy protection and unclonable cryptography in general. In particular, the identical-challenge variant (along with $i\mathcal{O}$) implies identical-challenge secure copy protection for various functionalities, including identical-challenge secure SDE and copy protection of point functions[2], as well as unclonable indistinguishable encryption. Unfortunately, [AB24] could only construct two variants of UPO based on two conjectures. In a later work, the requirement of conjectures was removed for the case of independent challenge distributions [ABH+26], but constructing the identical-challenge variant of UPO in the plain model still remains open, which brings us to the following natural question.

Can we achieve UPO with identical-challenge security and, more generally, correlated-challenge security? Moreover, can we achieve stronger generalizations of UPO that can help us obtain copy protection for more functionalities beyond what is currently known in the plain model?

1.1 Our Results

We note that correlation was the missing axis in prior constructions, as they primarily address independent challenges, whereas natural applications often give the two recipients correlated or identical challenges. Our broad conceptual contribution is that we provide a framework to handle correlation in challenge distributions that enables us to answer the questions asked above affirmatively.

Gold-standard definition for SDE. We introduce a new simple and natural security definition for SDE called *correlated-challenge security* (SDE-CORR) that better captures the security intuition and use cases for SDE, and also show that it implies all existing SDE security definitions.

Theorem 1 (Informal)**.** *SDE-CORR implies all seven prior security notions for SDE, including identical-challenge security.*

As discussed below, the relationships of the prior definitions among themselves is also unclear. In this work, we establish an almost complete categorization of the prior definitions, which in particular also shows their inadequacy: For example, the usual indistinguishability/CPA SDE security definition[3] does not even imply search security! Further, we show that independent-challenge CPA security does not imply identical-challenge CPA security, refuting a folklore conjecture and showing the inadequacy of previous techniques for proving identical-challenge CPA security. We refer the reader to Figure 1 for the complete picture.

We also achieve our gold-standard notion in the plain model, with same cryptographic assumptions as previous works.

Theorem 2 (Informal)**.** *Assuming indistinguishability obfuscation and one-way functions, there exists an SDE scheme satisfying SDE-CORR.*

[2] In fact, it implies a generalization of point functions called k-point functions.

[3] Here in indistinguishability/CPA SDE notion, we do not include non-operational CPA notions like Strong CPA anti-piracy [CLLZ21].

In particular, this is the first classical-ciphertext SDE construction satisfying identical-challenge security from standard cryptographic assumptions, resolving the question introduced by Georgiou and Zhandry and left open throughout subsequent work [GZ20, CLLZ21, AKL23, CGLZR23, KN23, CHV23, AB24, AKY25a, KY25].

Copy-Protecting general functionalities. For copy-protecting general functionalities, we follow the modular approach of Ananth and Behera [AB24], and achieve our feasibility results via unclonable puncturable obfuscation (UPO).

We first introduce a strengthened definition for UPO called correlated-challenge unclonable puncturable obfuscation (UPO) or simply correlated UPO, that implies both variants of UPO security definitions introduced by [AB24], i.e., identical-challenge security and independent-challenge security. Then, we show how to achieve it in the plain model.

Theorem 3 (Informal)**.** *Assuming indistinguishability obfuscation and hardness of LWE, there exists an unclonable puncturable obfuscation (UPO) scheme that satisfies correlated security for any sufficiently unpredictable challenge distribution.*

This gives, in particular, the first provably secure identical-challenge generalized UPO in the plain model and, through the UPO framework, the first identical-challenge-secure copy protection for puncturable functionalities, resolving the question left open by Çakan–Goyal and Ananth et al. [ÇG26, ABH+26], as well as new feasibility results for copy protection for various evasive function classes like point functions and compute-and-compare functions.

Our definition of correlated UPO security not only allows correlation between challenge distributions but also allows for auxiliary information leakage to the adversaries, which can be critical for various applications. Indeed, the reduction for the application to copy protection of compute-and-compare functions crucially relies on the auxiliary information leakage.

Applications. As discussed before, identical-challenge UPO security is related to many fundamental questions. Thus, our construction yields the first plain-model copy protection for point functions with the natural challenge distribution, resolving a question posed by Aaronson when he introduced quantum copy protection in 2009 [Aar09]. Our result holds for arbitrary high-min entropic distribution on point functions, and also generalizes to k-point functions, a generalization of point functions, considered in [AB24].

Beyond point functions, we obtain the first plain-model copy protection for compute-and-compare programs with a natural challenge distribution just as in point-functions[4], a more expressive class whose quantum copy protection was introduced and previously achieved only in the quantum random-oracle model by Coladangelo, Majenz, and Poremba [CMP20]. We note that for the application to copy protection of compute-and-compare functions, we rely on the strength of the new correlated UPO security, both in terms of the correlation in the challenge distribution and the auxiliary information leakage that the correlated UPO security definition allows for.

Theorem 4 (Informal)**.** *Assuming post-quantum indistinguishability obfuscation and quantum-hard LWE; point functions, k-point functions, and compute-and-compare programs admit plain-model copy protection with respect to natural identical challenge distributions, and arbitrary high min-entropic distribution on the respective function families. Moreover, for identical challenge*

[4]This natural challenge distribution is an identical challenge distribution, as for the case of point functions. We also note that though related, the proof of feasibility of copy protection for compute-and-compare functions does not follow directly from the feasibility of point functions.

distributions with uniform marginals for all three copy protection applications, post-quantum one-way functions suffice in place of LWE.

Beyond copy-protection for point functions, our feasibility result for correlated UPO extends the existing results of independent-challenge copy-protection for puncturable functionalities to correlated-challenge copy-protection. Moreover, in the identical-challenge setting, our results extend to stronger definitions of copy-protection including oracular security, introduced by [ABH+26].

Theorem 5 (Informal)**.** *Assuming post-quantum indistinguishability obfuscation and quantum-hard LWE, puncturable functionalities admit unpredictability and pseudorandomness-style copy protection with respect to* arbitrary correlated high min-entropic challenge distributions*. Moreover, for the case of* identical high min-entropic challenge distributions*, the oracular variant of the respective copy protection notion (i.e., unpredictability or pseudorandomness-style) that gives adversaries oracle access to the protected function holds.*

A new decisional monogamy game for coset state. We achieve our results through a new, simple-to-state *decisional* monogamy-of-entanglement guarantee we prove for coset states.

Theorem 6 (Informal; Decisional Coset Monogamy)**.** *Let $A \leq \mathbb{F}_2^{2n}$ be a uniformly random n-dimensional subspace, let $s, t \in \mathbb{F}_2^{2n}$ be uniform, and give an adversary one copy of*

$$|A_{s,t}\rangle := \frac{1}{\sqrt{|A|}} \sum_{a \in A} (-1)^{\langle a,t \rangle} |a + s\rangle .$$

Consider the game where an adversary splits this state between two noncommunicating adversaries, and a challenger samples uniform non-zero $u \in A$ and non-zero $v \in A^\perp$ sends it to both adversaries. The probability that both users correctly output the bit $e := \langle u, t\rangle \oplus \langle v, s\rangle$ is at most $1/2 + 2^{-\Omega(n)}$.

Moreover, assuming post-quantum indistinguishability obfuscation and subspace-hiding obfuscation, the corresponding computational version of the MOE guarantee holds even when the adversary receives the public obfuscated membership programs

$$M_0 = \mathsf{iO}(P_{A+s}) \qquad \text{and} \qquad M_1 = \mathsf{iO}(P_{A^\perp + t}),$$

together with correlated simulatable auxiliary information, similar to the computation variant considered in [ABH+26].

We emphasize that, unlike the search-style coset monogamy theorem of Culf and Vidick [CV22], our game is *decisional*, i.e., both recipients $\mathcal{B}$ and $\mathcal{C}$ output a *single-bit answer*, rather than a full vector from an affine subspace. Secondly, our monogamy theorem has truly *identical* challenges for both recipients, i.e., they receive the same challenge (as opposed to mildly correlated challenges) and output the *same single-bit answer*.

Furthermore, to the best of our knowledge, our monogamy of entanglement theorem is the first identical decisional monogamy of entanglement theorem with same-bit answers for an *unlearnable* state family, which is crucial for our purposes. Earlier decision-style monogamy of entanglement experiments with same bit answers (i.e., requires both recipients to output the same bit) including the celebrated work of [TFKW13] and the recent unclonable encryption result of [AS26], are based on *learnable* state families such as BB84 or Pauli states. Thus, they cannot be used for copy protection, even in simpler forms such as *point function copy protection*. Our result provides the missing combination of unlearnability and decision-form monogamy, unifying the study of monogamy games, unclonable encryption and copy protection. Our decisional monogamy of

entanglement theorem with same bit answers also significantly simplifies the copy-protection reductions: it avoids the rewinding and threshold-measurement techniques used for simultaneous extraction in previous works [Wat06, Zha20, ALL+21] and removes the need to canonicalize recovered subspace vectors, thereby giving a simpler alternative route to achieving copy protection, that does not require any extraction techniques. Hence, we believe our new monogamy of entanglement result will have independent applications.

1.2 Organization

Section 3.1 recalls notation and preliminaries. Sections 4.1 to 4.3 define the SDE notions and establish their hierarchy, after which Sections 5.1.1 and 5.1.2 define correlated-challenge UPO. The next part proves the information-theoretic and computational decisional coset-monogamy theorems. Finally, Section 8 constructs a correlated-challenge SDE, and Section 9 constructs a correlated secure UPO, and Section 10 derives the copy protection applications.

Acknowledgments

Alper Çakan was supported by the following grants of Vipul Goyal: NSF award 1916939, DARPA SIEVE program, a gift from Ripple, a DoE NETL award, a JP Morgan Faculty Fellowship, a PNC center for financial services innovation award, and a Cylab seed funding award.

AI Acknowledgment

Our information-theoretic MoE result and our SDE-CORR definition was proven with help of interaction with ChatGPT-5.6 Sol. Rest of the results were obtained directly by authors. ChatGPT-5.6 Sol was also used for drafting the initial versions of some text. The authors take full responsibility for the whole paper.

2 Technical Overview

In this section, we outline the key ideas and an overview of our proof techniques. We note that the constructions are largely the same as considered in previous works, namely, the construction of [KY25] for the single-decryptor encryption (SDE) construction, and the construction of [ABH+26, ÇG26] for the unclonable puncturable obfuscation (UPO) construction. We very briefly recall the definitions of SDE and UPO in the literature. An SDE scheme [CLLZ21, KY25] is a public key encryption scheme with classical ciphers and public keys but equipped with a quantum decryption key that can decrypt any cipher generated using the public key, can be correctly decrypted using the quantum decryption key. In the CPA security experiment, there is a triplet of adversaries $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ such that $\mathcal{A}$ receives a quantum decryption key along with the public key from the challenger, and outputs two pairs of challenge messages, one for $\mathcal{B}$ and one for $\mathcal{C}$, and then sends two (potentially entangled) quantum registers, one to $\mathcal{B}$ and the other to $\mathcal{C}$ who cannot communicate among themselves. Then, the challenger sends each of $\mathcal{B}$ and $\mathcal{C}$ the encryption of one of the two challenge messages chosen according to a randomly sampled challenge bit for each recipient, and then $\mathcal{B}$ and $\mathcal{C}$ need to guess their respective challenge bits. In the identical challenge setting, $\mathcal{A}$ only chooses a single pair of challenge messages, and the challenger uses the same uniformly random bit for both $\mathcal{B}$ and $\mathcal{C}$, who receive the exact same cipher text, and needs to predict their common bit, whereas in the independent challenge setting, the challenge bits for $\mathcal{B}$ and $\mathcal{C}$ are sampled uniformly at random independently. In this work, we also consider a generalization of the independent challenge setting, where we allow the challenge bits of $\mathcal{B}$ and $\mathcal{C}$ to be arbitrarily correlated as long as the marginal distribution for each of the two challenge bits is uniform.

An unclonable puncturable obfuscation (UPO) [AB24] is a tuple of QPT algorithms $\mathsf{Obf}, \mathsf{Eval}$ where Obf compiles a classical circuit C into a quantum program ρ_C, and Eval takes as input a quantum program ρ_C and an input x, and outputs $C(x)$. In the identical challenge UPO security experiment, again there are three adversaries $(\mathcal{A}, \mathcal{B}, \mathcal{C})$. $\mathcal{A}$ selects a circuit C to the challenger, who then samples challenge bit b, and a single challenge input point x sampled from a high min-entropy distribution, and if the bit $b = 0$, the challenger runs Obf on C else if $b = 1$, it first punctures the circuit at the point x using a puncturing algorithm equipped with the circuit class[5], and sends the program to $\mathcal{A}$ who then splits the state in to two quantum registers and sends one to $\mathcal{B}$ and one to $\mathcal{C}$, after which the puncture point x is revealed to both $\mathcal{B}$ and $\mathcal{C}$ and then both $\mathcal{B}$ and $\mathcal{C}$ need to output the bit b. In the independent challenge settings, the only difference is that there are two challenge points sampled independently for $\mathcal{B}$ and $\mathcal{C}$, whereas in general correlated challenge setting, we allow the two challenge points to be arbitrarily correlated.

Existing approaches and their limitations. For a subspace $A \leq \mathbb{F}_2^q$ of dimension $q/2$, and vector shifts $s, t \in \mathbb{F}_2^q$, coset states refer to states of the form

$$|A_{s,t}\rangle := \frac{1}{\sqrt{|A|}} \sum_{a \in A} (-1)^{\langle a,t \rangle} |a + s\rangle ,$$

and we refer to $H^{\otimes q} |A_{s,t}\rangle = \left|A^{\perp}_{t,s}\right\rangle := \frac{1}{\sqrt{|A^{\perp}|}} \sum_{a \in A^{\perp}} (-1)^{\langle a,s \rangle} |a + t\rangle$, as the corresponding dual coset for the dual subspace $A^{\perp}$. Coset states are a natural candidate of state family to build SDE, and copy protection in general as well as UPO, because they have two desirable yet contrasting

[5] By puncturing a circuit at a point x, we just mean outputting a circuit that outputs the true value $C(x')$ on all points $x' \neq x$ and outputs $\perp$ on input x. Clearly, any circuit class can be equipped with at least the canonical puncturing algorithm that given a circuit C and point x produces such a circuit by hardcoding C and x into a new circuit C_x.

properties, namely, they can be verified publicly, but they still satisfy a strong form of unclonability, colloquially referred to as monogamy of entanglement (MOE) [CV22, CLLZ21]. In particular, the monogamy of entanglement guarantees that, given a random coset state, $\mathcal{A}$ cannot split the state into two registers and send them to $\mathcal{B}$ and $\mathcal{C}$ then, given the subspace A, $\mathcal{B}$ and $\mathcal{C}$ cannot recover the vector shifts s and t, respectively.

How coset monogamy of entanglement was used in previous works. The starting point in previous works is an information-theoretic search guarantee: after $\mathcal{A}$ splits a coset state, $\mathcal{B}$ and $\mathcal{C}$ cannot both output suitable elements of the primal coset $A + s$ and the dual coset $A^{\perp} + t$. Here "search" means that they must output complete coset elements, rather than one-bit answers. Coladangelo, Liu, Liu, and Zhandry [CLLZ21] proved a computational version in which the cosets are represented by obfuscated membership programs, namely, programs that check whether an input belongs to the primal or dual coset without revealing the cosets themselves. Their proof uses subspace-hiding obfuscation, which hides certain changes to the subspaces represented by these programs, together with information-theoretic properties of coset states.

The SDE and UPO security games ask $\mathcal{B}$ and $\mathcal{C}$ to predict bits. Previous proofs therefore need a search-to-decision reduction connecting these one-bit predictions to the search guarantee above. After a sequence of hybrids using $i\mathcal{O}$ and other cryptographic properties, these proofs canonicalize the accepting inputs: among the many coset elements accepted by an $i\mathcal{O}$-obfuscated membership program, they specify one value to be extracted as the search output. The proof of [CLLZ21] performs this extraction using threshold implementations and compute-and-compare obfuscation. A threshold implementation is a quantum measurement that tests whether an algorithm's success probability is above a chosen threshold, while a compute-and-compare program accepts when a hidden computation equals a specified value. The proofs of [ÇG26, ABH+26] instead use the simultaneous Goldreich–Levin theorem of Ananth, Kaleoglu, and Yuen [AKY25a], which turns sufficiently good inner-product predictors on both sides into the corresponding hidden values.

Why the previous techniques do not handle identical challenges. Extraction is particularly difficult when the two registers may be entangled. Extracting a coset element from $\mathcal{B}$'s register may disturb $\mathcal{C}$'s register and destroy the guarantee needed on his side. Previous works handle this issue in the independent-challenge setting by critically using the independence of the two challenges via threshold implementations of local tests. However, these techniques do not work when $\mathcal{B}$ and $\mathcal{C}$ receive identical challenges.

One approach used in previous works to avoid search-to-decision reduction and hence extraction was to use MOE guarantees where $\mathcal{B}$ and $\mathcal{C}$ only need to output a single-bit answer. However, the only MOE experiment with such a property, which was proved by Kitagawa and Yamakawa [KY25], is not an *identical-challenge* MOE guarantee, as the basis for it is the simultaneous Goldreich-Levin theorem [AKY25b], and hence gives the $\mathcal{B}$ and $\mathcal{C}$ independently sampled random strings as part of their challenges. Thus, the existing coset monogamy of entanglement guarantees are either search-based [CLLZ21, CV22] or are not with respect to identical challenges, whereas to circumvent the extraction barrier, we need a decisional guarantee in which $\mathcal{B}$ and $\mathcal{C}$ receive the same challenge, and both must output the same single bit.

2.1 Our Solution: Identical-Challenge Decisional Coset Monogamy of Entanglement

We prove the following truly identical-challenge decisional coset monogamy of entanglement theorem.

Identical-Challenge Decisional Coset Monogamy of Entanglement

1. The challenger samples a uniformly random $q/2$-dimensional subspace $A \leq \mathbb{F}_2^q$, uniform shifts $s, t \in \mathbb{F}_2^q$, and gives $|A_{s,t}\rangle$ to $\mathcal{A}$.
2. $\mathcal{A}$ splits the state into two registers and sends one to $\mathcal{B}$ and the other to $\mathcal{C}$.
3. The challenger samples $u \xleftarrow{\$} A \setminus \{0\}$ and $v \xleftarrow{\$} A^\perp \setminus \{0\}$, and gives the same pair (u, v) to both $\mathcal{B}$ and $\mathcal{C}$.
4. $\mathcal{B}$ and $\mathcal{C}$ output bits $e_\mathcal{B}$ and $e_\mathcal{C}$, respectively. They win if $e_\mathcal{B} = e_\mathcal{C} = e := \langle u, t\rangle \oplus \langle v, s\rangle$.

We will show that for every adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$,

$$\Pr[e_\mathcal{B} = e_\mathcal{C} = e] \leq \frac{1}{2} + \frac{1}{2}\sqrt{\frac{2^q}{(2^q-1)(2^{q-1}-1)}} = \frac{1}{2} + 2^{-\Omega(q)}. \tag{1}$$

Remark 1. *A nice property of our decisional coset MOE guarantee is that the final answer given by $e := \langle u, t\rangle \oplus \langle v, s\rangle$ remains unchanged even if we replace s and t with arbitrary coset elements in $A + s$ and $A^\perp + t$. This is especially useful while proving the security of SDE and UPO, since now we do not need the values of s and t, or their canonical representatives, to calculate the required parity for the decisional coset MOE game. Hence, the reduction from SDE or UPO to the decisional coset MOE does not do any form of extraction of canonical representatives for the cosets.*

(For the formal game and theorem, see Definition 38 and Theorem 19.)

2.1.1 Information-Theoretic Security Proof

We first express the game above as a one-bit unclonable-encryption scheme based on coset states. For $a \in \mathbb{F}_2^q$, define the Pauli operators X^a and Z^a by $X^a |x\rangle := |x \oplus a\rangle$ and $Z^a |x\rangle := (-1)^{\langle a,x\rangle} |x\rangle$.

Coset-UE

1. GenKey: Sample a uniformly random nonzero orthogonal pair $k = (u, v)$, meaning that $u, v \in \mathbb{F}_2^q \setminus \{0\}$ and $\langle u, v\rangle = 0$.
2. $\mathsf{Enc}(k, e)$: Sample a uniformly random $q/2$-dimensional subspace A satisfying $u \in A$ and $v \in A^\perp$. Sample $s, t \in \mathbb{F}_2^q$ uniformly subject to $\langle u, t\rangle \oplus \langle v, s\rangle = e$, and output $|A_{s,t}\rangle$.
3. $\mathsf{Dec}(k, \rho)$: Measure $X^u Z^v$ on ρ and output the bit e corresponding to the measured eigenvalue $(-1)^e$.
4. Security: Give an encryption of a uniform bit e to $\mathcal{A}$. After $\mathcal{A}$ sends one register to each recipient, reveal the same key (u, v) to both. They win only if both output e.

Every nonzero orthogonal pair (u, v) belongs to the same number of subspaces A satisfying $u \in A$ and $v \in A^\perp$, and for fixed A, u, v, the bit $\langle u, t\rangle \oplus \langle v, s\rangle$ is uniform over s, t. Consequently, sampling the coset state first and (u, v) afterwards, as in the monogamy game, gives the same joint distribution as sampling (u, v) first and then encrypting, as in Coset-UE. (For the formal equivalence, see Lemma 8.)

Conditional state of the ciphertext given the message and the key. Let $|A\rangle := |A_{0,0}\rangle$. Directly from the definition of a coset state and the Pauli commutation rule, we have

$$X^x Z^z |A\rangle = |A_{x,z}\rangle \quad \text{and} \quad Z^v X^s = (-1)^{\langle v,s\rangle} X^s Z^v. \tag{2}$$

Since $|A_{s,t}\rangle = X^s Z^t |A\rangle$, these identities give

$$X^u Z^v |A_{s,t}\rangle = (-1)^{\langle v,s\rangle} |A_{s+u,t+v}\rangle . \tag{3}$$

We now use $u \in A$ and $v \in A^\perp$. Expanding the state on the right and changing the summation variable from a to $a+u$ gives

$$|A_{s+u,t+v}\rangle = \frac{1}{\sqrt{|A|}} \sum_{a\in A} (-1)^{\langle a,t+v\rangle} |a+s+u\rangle = (-1)^{\langle u,t\rangle} |A_{s,t}\rangle . \tag{4}$$

The condition $v \in A^\perp$ is used here to remove $\langle a, v\rangle$, while $u \in A$ ensures that $a \mapsto a+u$ is a change of variable within A. Combining the last two equations yields

$$X^u Z^v |A_{s,t}\rangle = (-1)^{\langle u,t\rangle \oplus \langle v,s\rangle} |A_{s,t}\rangle . \tag{5}$$

Thus a Coset-UE encryption of e is an eigenvector of $X^u Z^v$ with eigenvalue $(-1)^e$.

Next, fix the key $k = (u, v)$ and encrypted bit e. The equation above shows that every possible ciphertext lies in the $(-1)^e$-eigenspace of $X^u Z^v$. We next explain why the ciphertext is in fact uniformly distributed over this entire eigenspace.

Fix any subspace A satisfying $u \in A$ and $v \in A^\perp$. Two shifts s that differ by an element of A describe the same shifted subspace $A+s$, while all other choices give disjoint sets of computational-basis vectors. For a fixed shifted subspace, choices of t that differ by an element of $A^\perp$ give the same relative signs, while all other choices give orthogonal sign patterns over A. The resulting 2^q coset states therefore form an orthonormal basis of the full 2^q-dimensional space. Since $X^u Z^v$ is a non-identity Pauli operator, each of its two eigenspaces has dimension 2^{q-1}. Exactly half of this coset basis has eigenvalue $(-1)^e$, so the allowed ciphertexts form an orthonormal basis of that eigenspace. Moreover, the uniform choices of s, t make the mixture uniform on this basis.

Writing $N := 2^q$ and letting $\Pi_{k,e}$ be the projector onto the $(-1)^e$-eigenspace, we obtain

$$\Pi_{k,e} = \frac{I + (-1)^e X^u Z^v}{2} \quad \text{and} \quad \rho_{k,e} = \frac{2}{N}\Pi_{k,e} = \frac{1}{N}\big(I + (-1)^e X^u Z^v\big). \tag{6}$$

Crucially, the final expression does not depend on A. Thus, averaging over the randomly chosen subspaces satisfying $u \in A$ and $v \in A^\perp$ leaves the same state. This is our main technical observation: conditioned on the key and message bit, a Coset-UE ciphertext is simply the maximally mixed state on one eigenspace of $X^u Z^v$. (For the complete calculation, see Claim 1.)

Using the argument of Ananth and Sahai [AS26]. The form of the conditional ciphertext in (6) is exactly the structure used in their proof of unclonable encryption. We prove a generalization for any family of Hermitian operators $\{\Theta_k\}_k$ on an N-dimensional space such that each Θ_k has eigenvalues $+1$ and -1 with equally large eigenspaces, and $\mathrm{Tr}(\Theta_k \Theta_{k'}) = 0$ for distinct keys. If a key is sampled from a distribution P and the encryption of e is $N^{-1}(I + (-1)^e \Theta_k)$, then the probability that both recipients recover e is at most

$$\frac{1}{2} + \frac{1}{2}\sqrt{N \sum_k P(k)^2}. \tag{7}$$

The quantity $\sum_k P(k)^2$ is the collision probability of the distribution on the key space, namely, the probability that two independently sampled keys are equal. The proof uses the same argument as [AS26]; it depends only on the fact that the ciphertext state is the normalized projection onto the eigenspace, that there are only $1, -1$ eigenvalues, and that both eigenspaces have equal dimensions, which follows from the abstractions mentioned above. It is easy to see that the unclonable-encryption scheme of [AS26] satisfies these conditions. Coset-UE also satisfies them because $\mathrm{Tr}\Big((X^u Z^v)(X^{u'} Z^{v'})\Big) = 0$ for distinct Pauli operators. Finally, substituting the uniform distribution over nonzero orthogonal pairs gives the required bound in (1), proving our information-theoretic decisional coset monogamy of entanglement theorem. (For the general statement and proof, see Lemma 9.)

2.1.2 Computational Decisional Hidden-Coset Monogamy

The eigenspace identity above averages over the hidden subspace and shifts. In the security experiment for SDE and UPO, the adversary also receives membership programs and auxiliary strings correlated with these hidden values. We cannot replace the coset states by their conditional average without first handling the correlated information. Therefore, we follow the modular template of [CLLZ21, ABH+26] to first present an additional computational version of our decisional coset MOE guarantee that would help us prove the security of both SDE and UPO.

The computational variant gives $\mathcal{A}$ the coset state together with obfuscated membership programs for $A+s$ and $A^\perp+t$; these programs test membership while hiding A, s, t. It also allows $\mathcal{A}$ to receive classical information before splitting the state, as long as this information can be generated without the later pair (u, v), and it allows separate classical information to be given to $\mathcal{B}$ and $\mathcal{C}$ after (u, v) is sampled. Even with this information, no efficient adversary can make both parties recover $\langle u, t\rangle \oplus \langle v, s\rangle$ with probability noticeably greater than $1/2$. (For the formal computational game and theorem, see Definition 39 and Theorem 21.)

The proof mainly uses the previous computational coset monogamy of entanglement techniques of [ABH+26, CLLZ21]. In particular, subspace-hiding obfuscation lets the proof change the subspaces represented by the public membership programs without an efficient adversary noticing, and the algebraic properties of coset states used in those works remove the remaining dependence on the original hidden coset. The classical information given before the split is then replaced by a separately generated copy with the same distribution that does not use (u, v). After these changes, the remaining experiment is the information-theoretic identical challenge-decisional game proved above. Thus, the proof of the computational MOE theorem follows by combining the earlier computational techniques with our new information-theoretic theorem.

2.2 Correlated-Challenge Single-Decryptor Encryption

Independent challenges do not capture the setting where a ciphertext is generated once and then copied to several users, or when the same or correlated messages are encrypted separately to get potentially different challenge distributions. We therefore want a single definition that includes independent encryption executions, literal copies of one ciphertext, and arbitrary correlations between the two challenge bits, subject to uniform marginals. We define such a notion of SDE security as follows.

Correlated-challenge SDE security. The experiment begins as in the usual SDE game: the challenger gives $(\mathsf{pk}, \mathsf{sk})$ to $\mathcal{A}$, who prepares one register for $\mathcal{B}$ and one for $\mathcal{C}$. In addition, $\mathcal{A}$ chooses *polynomially many* challenge-message pairs. The use of several pairs means that one challenge bit

chooses the left or right message in every pair; it does not mean that a new challenge bit is sampled for each pair.

The experiment has two modes. In separate mode, $\mathcal{A}$ may choose different message pairs for $\mathcal{B}$ and $\mathcal{C}$. The challenger samples $(b_{\mathcal{B}}, b_{\mathcal{C}})$ from any efficiently samplable joint distribution such that the marginal distribution on each of the two individual bits is uniform, and independently encrypts one of the two messages from the polynomially many pairs of challenge messages according to the bit b_X, and gives the resulting ciphertexts to recipient X. Note that the two bits may be independent, or arbitrary correlated. In identical mode, $\mathcal{A}$ chooses the same message pairs for both $\mathcal{B}$ and $\mathcal{C}$, and the challenger samples one uniform bit b, generates one ciphertext for each message selected by b, and gives an exact copy of every resulting ciphertext to both recipients. The scheme is secure if no efficient $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ makes both recipients recover their respective challenge bits with probability noticeably greater than $1/2$. When there is only one challenge-message pair, the two modes reduce to the usual independent-CPA and identical-CPA experiments. (For the formal sampler and security game, see Definitions 21 and 22.)

Construction. We use the coset-based SDE construction of [KY25]. Key generation samples A, s, t, publishes obfuscated programs testing membership in $A + s$ and $A^{\perp} + t$, and uses $|A_{s,t}\rangle$ as the quantum decryption key. To encrypt a message μ, the encryption algorithm chooses a uniformly random string r, called the pad, and obfuscates a program with the following behavior: on an input from $A + s$, it returns $\mu \oplus r$, while on an input from $A^{\perp} + t$, it returns r. An honest decryptor evaluates the first behavior using the coset state, applies Hadamard gates to obtain the dual coset state, evaluates the second behavior, and sum the two answers to recover μ. Both evaluations are reversed, so the quantum decryption key is restored.

Security proof: reduction to computational decisional coset monogamy of entanglement. We follow the proof template of [KY25], but now in the setting of correlated ciphertext distribution. In order to handle arbitrary correlation among challenge bits $b_{\mathcal{B}}$ and $b_{\mathcal{C}}$, we observe that it is possible to rewrite arbitrarily correlated challenge bits using one identical hidden bit, without changing the overall distribution. Indeed, two bits that are individually uniform can be sampled by first deciding whether they should be equal or opposite and then sampling their common random value. Let Δ record this first choice: $\Delta = 0$ means that the bits are equal and $\Delta = 1$ means that they are opposite. Their joint distribution can then be written as $(g, g \oplus \Delta)$ for a uniform bit g.

The computational coset monogamy game provides the hidden bit $e = \langle u, t\rangle \oplus \langle v, s\rangle$. We sample another uniform bit ζ and set

$$b_{\mathcal{B}} = e \oplus \zeta \quad \text{and} \quad b_{\mathcal{C}} = e \oplus \zeta \oplus \Delta. \tag{8}$$

Because $e \oplus \zeta$ is uniform, these two bits have the desired joint distribution. Moreover, correct answers from $\mathcal{B}$ and $\mathcal{C}$ immediately give two answers to the same monogamy challenge: Add $\mathcal{B}$'s answer with ζ, and add $\mathcal{C}$'s answer with $\zeta \oplus \Delta$. If both recipients answered correctly, both resulting bits equal e.

The remaining question is how to create the corresponding ciphertexts without knowing the shifts s and t. We rewrite the ciphertext program so that its value on $A + s$ depends only on $\langle v, s\rangle$, while its value on $A^{\perp} + t$ depends only on $\langle u, t\rangle$. These two bits can be computed from the program's input: $\langle v, w\rangle = \langle v, s\rangle$ for $w \in A + s$, and $\langle u, w\rangle = \langle u, t\rangle$ for $w \in A^{\perp} + t$. The rewritten program agrees with an honestly generated ciphertext program on every input, so $i\mathcal{O}$ allows us to replace one by the other. This is the same programming idea used in the proof of Kitagawa and

Yamakawa [KY25]; the new point is that the equal-or-opposite representation above makes it work for every allowed correlation between the two challenge bits.

Thus, if both SDE recipients recover their challenge bits, the two summing corrections described above make both of them recover e, contradicting computational decisional coset monogamy of entanglement. In identical mode, we use the same message pair, pad, program, and obfuscation randomness for both recipients and copy the resulting classical ciphertext. The proof therefore preserves the fact that the two recipients receive exactly the same ciphertext. (For the formal construction and reduction, see Section 8 and Theorem 22.)

2.3 Copy Protection from Correlated UPO

Given a UPO scheme $(\mathsf{Obf}, \mathsf{Eval})$, we protect a circuit C by running $\mathsf{Obf}(C)$ and evaluate it using Eval. The security proof changes the value of C at the challenge points. UPO ensures that two recipients cannot distinguish the protected original circuit from the protected modified circuit, while the classical assumptions on C ensure that its original values at those points remain hidden after the change. This follows the framework of [AB24, ABH+26]. The new point is that our proofs preserve the correlation between the two challenge points and modify the circuit only once when those points are equal.

Point functions and k-point functions. A point function P_y outputs one only at the marked point y and zero everywhere else. If an independently sampled challenge point has high min-entropy, then it equals y only with negligible probability. In this independent case, we can therefore replace the required answer by zero and then use UPO security to replace the protected point function by the protected all-zero circuit. This extends the proof of [AB24] from uniformly random points to arbitrary high-min-entropy distributions. The same argument applies to functions that output one on k marked points, under the corresponding sampling and entropy requirements. (For the formal statements, see Theorem 25 and Corollary 5.)

Compute-and-compare functions and auxiliary information. A compute-and-compare function has the form $\mathsf{CC}_{f,y}(x) = [f(x) = y]$. The natural construction publishes f, protects the point function P_y, and evaluates it on $f(x)$, as in [CMP24]. Its security does not follow immediately from point-function copy protection: the public function f may be correlated with y, and the recipients are challenged with a preimage x, not directly with the point $f(x)$. A reduction must therefore give f to $\mathcal{A}$ before she sends out the two registers, but give x to $\mathcal{B}$ and $\mathcal{C}$ only afterwards.

Correlated UPO allows precisely this order of information. We give f to $\mathcal{A}$ first, use $z = f(x)$ as the common UPO challenge point, and give the same preimage x to both recipients later. The required entropy condition says that z remains difficult to guess even after f is known. An attack on the resulting compute-and-compare copy protection would therefore give an attack on correlated UPO. This is why the auxiliary-information feature of our UPO definition is needed for this application. (For the formal result, see Theorem 26.)

General function classes with correlated challenges. The transformations of [AB24, ABH+26] also extend to general functions satisfying the required classical security properties. For unpredictability-style security, where both recipients must compute the correct function values, we change the circuit at the distinct challenge points. For pseudorandomness-style security, where each recipient must distinguish the real value from a random one, we change only the values selected to be random. Throughout both arguments, the two challenge points remain jointly distributed. If they are equal,

the circuit is changed only once, following the same convention as in the correlated UPO experiment. (For the formal transformations, see Theorems 27 and 29.)

Oracle access with an identical challenge. Here $\mathcal{A}$ receives the protected program and oracle access to the original function C. After she sends one register to each recipient, the challenger samples a high-min-entropy point x and gives the same x to both recipients. In addition, $\mathcal{B}$ receives oracle access to $C^{\setminus x}$, which agrees with C everywhere except that it returns $\perp$ at x, and $\mathcal{C}$ receives oracle access to this same function $C^{\setminus x}$.

The only additional issue is $\mathcal{A}$'s earlier access to C. The point x is sampled only after she sends the two registers and is difficult to guess, so her polynomially many oracle queries are very unlikely to involve x. A standard quantum-query argument therefore lets us replace her oracle by $C^{\setminus x}$ with only a negligible change. Since $\mathcal{B}$ and $\mathcal{C}$ already receive the same point and the same punctured oracle, the rest of the proof follows directly from correlated UPO. (For the formal results, see Theorems 28 and 30.)

3 Preliminaries

3.1 Notation and Conventions

We refer the reader to [Gol01, Gol04] for a preliminary on cryptography and to [NC10] for a preliminary on quantum information and computation. We follow the notations and conventions commonly used in (quantum) cryptography, theoretical computer science and mathematics. In this section, we non-exhaustively highlight some of them. All of our conventions are implicitly followed unless otherwise specified, e.g. our cryptographic assumptions are always post-quantum even though most of the time we will not write this explicitly, and if there is a rare occasion where we assume classical security only, we will specify this explicitly.

Functions For a function $f : \mathbb{N}^+ \to \mathbb{N}^+$, we will say that it is an efficiently computable function if it is polynomially bounded, that is, if there is some $n_0 \in \mathbb{N}^+$ and $c \in \mathbb{R}^+$ such that $f(n) \leq n^c$ for all $n > n_0$. Sometimes we will say polynomial to also mean polynomially bounded instead (e.g. we will say polynomial to refer to $O(\mathsf{poly}(\lambda))$, which includes, say, $\sqrt{n}$) - the precise meaning will be clear from context. $f(\lambda)$ is said to be negligible if for any $c \in \mathbb{N}^+$, there is $n_0 \in \mathbb{N}^+$ such that $f(\lambda) \leq \frac{1}{n^c}$ for all $n \geq n_0$. $f(\lambda)$ is said to superlogarithmic if $f = \omega(\log(\lambda))$ or equivalently, if $2^{-f(\lambda)}$ is negligible.

When we describe a function, such as a polynomial, if no input is mentioned, it will implicitly be the security parameter λ.

Classical and Quantum Information and Oracles Variables written in lower-case, such as pk, sk or pk, sk, are classical objects. Variables written in the font $\mathscr{sk}$ are quantum variables. We write R to mean a quantum register, which keeps a quantum state that will evolve when the register is acted on, and it can be entangled with other registers.

We write $\mathsf{R} \leftarrow \rho$ to mean that the register R is initialized with a sample from the quantum distribution (i.e. mixed state) ρ. For simplicity, we will usually write non-normalized versions of the quantum states, however, they are always implicitly normalized.

For mixed states ρ and σ, we write $\|\rho - \sigma\|_{Tr} := \frac{1}{2}\|\rho - \sigma\|_1$ for their trace distance.

We write $\mathcal{O}$ to denote a classical oracle, and write $\mathcal{A}^{\mathcal{O}}$ to denote an algorithm that has quantum query-access to $\mathcal{O}$. Similarly, for unitary oracles we write U and $\mathcal{A}^U$. All adversaries always have quantum access to all of the oracles.

Distributions When S is a set, we write $x \leftarrow S$ to mean that x is sampled uniformly at random from S. When $\mathcal{D}$ is a distribution, we write $x \leftarrow \mathcal{D}$ to mean that x is sampled from $\mathcal{D}$. Finally, we write $x \leftarrow \mathcal{B}(\mathsf{R})$ or $x \leftarrow \mathcal{B}(a)$ to mean that x is a sample as output by the (quantum or classical randomized) algorithm $\mathcal{B}$ run on the quantum input register R or classical input a. We use similar notation for quantum registers, e.g., $\mathsf{R} \leftarrow \mathcal{B}(a)$.

Adversaries and Constructions/Algorithms We write *QPT* to mean *quantum polynomial time* and *PPT* to mean *probabilistic (classical) polynomial time*. We say query-bounded to mean that the algorithm (e.g. construction, adversary) makes polynomially many queries to the oracle, but it can possibly take unbounded time outside the queries. All of these are polynomial in the security parameter λ. Note that for QPT and PPT, this also means that the inputs to these algorithms are also of polynomial size in λ.

We say *efficient* to mean QPT or PPT (will be clear from context) in the plain model, or query-bounded in the oracle model. For constructions, efficient means QPT or PPT, even in the oracle model (though the oracles themselves might not be efficiently sampleable or implementable).

In the plain model, adversaries will be implicitly QPT. In the oracle security model, adversaries will be implicitly *query-bounded*. An admissible adversary means an adversary that conforms to the (implicit) requirements of the security game (e.g. polynomial time or query bounded, has the correct interaction model and has correct input-output format and so on). All of our cryptographical assumptions are post-quantum, e.g., *one-way functions* means *post-quantum secure one-way functions*. Adversaries are stateful uniform QPT.

When we say that a primitive is subexponentially secure, we mean that any polynomial time adversary has subexponentially small advantage. Sometimes we require security against subexponential time adversaries, in which case we will specify this explicitly.

All algorithms of a cryptographic scheme are uniform (stateless) PPT or QPT (will be clear from context).

Security Games A security game is a binary random variable which is the outcome of an experiment between a *challenger* and an adversary. For a security game, we say that the adversary has won if the experiment output is 1, and otherwise we say that the adversary has lost. When we say that an adversary has negligible advantage, we mean that the probability of the adversary winning (or distinguishing, depending on context) is upper bounded by a negligible function of λ.

Domains We write $\mathcal{M}$ to denote the message space for various cryptographic primitives, and it will be equal to $\{0,1\}^{p(\lambda)}$ for some polynomial $p(\cdot)$ (will be clear from context).

3.2 Pseudorandom Functions

In this section, we recall pseudorandom functions (PRF) and puncturable PRFs.

Definition 1 (Pseudorandom Functions)**.** *A pseudorandom function (PRF) scheme* $\mathbf{PRF}$ *is a family of functions* $\{F : \{0,1\}^{c(\lambda)} \times \{0,1\}^{m(\lambda)} \to \{0,1\}^{n(\lambda)}\}_{\lambda\in\mathbb{N}^+}$ *along with the following efficient algorithms.*

- $\mathbf{PRF}.\mathsf{Gen}(1^\lambda)$: *Takes in the unary representation of the security parameter and outputs a key* K *in* $\{0,1\}^{c(\lambda)}$*.*
- $\mathbf{PRF}.\mathsf{Eval}(K,x)$: *Takes in a key* $K \in \{0,1\}^{c(\lambda)}$ *and an input* $x \in \{0,1\}^{m(\lambda)}$*, outputs an evaluation of the PRF in* $\{0,1\}^{n(\lambda)}$*.*

We require the following.

Correctness.

$$\Pr\left[\forall x\ \mathbf{PRF}.\mathsf{Eval}(K,x) = F(K,x) :\ \ K \leftarrow \mathbf{PRF}.\mathsf{Gen}(1^\lambda)\ \right] = 1.$$

Security *We require that for any QPT adversary* $\mathcal{A}$*,*

$$\left|\Pr_{H\leftarrow \mathrm{H}}[\mathcal{A}^H(1^\lambda)=1] - \Pr_{K\leftarrow \mathbf{PRF}.\mathsf{Gen}(1^\lambda)}[\mathcal{A}^{F(K,\cdot)}(1^\lambda)=1]\right| \leq \mathsf{negl}(\lambda).$$

where H *is the uniform distribution over the functions* $\{0,1\}^{m(\lambda)} \to \{0,1\}^{n(\lambda)}$*.*

We will usually write $F_K(x)$ or $F(K,x)$ in algorithms and this will implicitly mean $\mathbf{PRF}.\mathsf{Eval}(K,x)$.

Definition 2 (Puncturable PRF)**.** *A puncturable PRF scheme* $\mathbf{PRF}$ *is a PRF scheme together with the following efficient algorithm.*

- $K_{\{\mathcal{S}\}} \leftarrow \mathbf{PRF}.\mathsf{Puncture}(K,\mathcal{S})$ *takes a key* $K \in \{0,1\}^{c(\lambda)}$ *and a polynomial-size set* $\mathcal{S} \subseteq \{0,1\}^{m(\lambda)}$*, and outputs a punctured key* $K_{\{\mathcal{S}\}}$*.*

We require the following properties.

Punctured correctness. *For every polynomial-size set* $\mathcal{S} \subseteq \{0,1\}^{m(\lambda)}$ *and every* $x \in \{0,1\}^{m(\lambda)} \setminus \mathcal{S}$*,*

$$\Pr\Big[F_{K_{\{\mathcal{S}\}}}(x) = F_K(x) : K \leftarrow \mathbf{PRF}.\mathsf{Gen}(1^\lambda),\ K_{\{\mathcal{S}\}} \leftarrow \mathbf{PRF}.\mathsf{Puncture}(K,\mathcal{S})\Big] = 1.$$

Pseudorandom at punctured point. *For every polynomial-size set* $\mathcal{S} \subseteq \{0,1\}^{m(\lambda)}$ *and every QPT distinguisher* $\mathcal{A}$*,*

$$\Big|\Pr\Big[\mathcal{A}\Big(F_{K_{\{\mathcal{S}\}}}, \{F_K(x)\}_{x\in\mathcal{S}}\Big) = 1\Big] - \Pr\Big[\mathcal{A}\Big(F_{K_{\{\mathcal{S}\}}}, (\mathcal{U}_{n(\lambda)})^{|\mathcal{S}|}\Big) = 1\Big]\Big| \leq \mathsf{negl}(\lambda),$$

where in both probabilities

$$K \leftarrow \mathbf{PRF}.\mathsf{Gen}(1^\lambda), \qquad K_{\{\mathcal{S}\}} \leftarrow \mathbf{PRF}.\mathsf{Puncture}(K,\mathcal{S}),$$

and $\mathcal{U}_{n(\lambda)}$ *is the uniform distribution over* $\{0,1\}^{n(\lambda)}$*.*

If $\mathcal{S} = \{x_1^\star, \ldots, x_q^\star\}$*, we write* $K_{\{x_1^\star,\ldots,x_q^\star\}}$ *for* $K_{\{\mathcal{S}\}}$*. For brevity, we sometimes write* $K_{\mathcal{S}}$ *as simply* $K_{\{\mathcal{S}\}}$*.*

It is easy to see that any puncturable PRF scheme that satisfies puncturing security also satisfies usual PRF security.

Theorem 7 ([SW14, Zha12])**.** *Let* $n(\cdot), m(\cdot)$ *be polynomially bounded.*

- *If (post-quantum) one-way functions exist, then there exists a (post-quantum) puncturable PRF with input space* $\{0,1\}^{m(\lambda)}$ *and output space* $\{0,1\}^{n(\lambda)}$*.*
- *If subexponentially-secure (post-quantum) one-way functions exist, then for any* $c > 0$*, there exists a (post-quantum)* $2^{-\lambda^c}$*-secure*[6] *puncturable PRF with input space* $\{0,1\}^{m(\lambda)}$ *and output space* $\{0,1\}^{n(\lambda)}$*.*

3.3 Digital Signature Schemes

In this section we recall the definition of signatures schemes.

Definition 3. *A digital signature scheme with message space* $\mathcal{M}$ *consists of the following efficient algorithms that satisfy the correctness and security guarantees below.*

[6]While the original results are for negligible security against polynomial time adversaries, it is easy to see that they carry over to subexponential security. Further, by scaling the security parameter by a polynomial and simple input/output conversions, subexponentially secure (for any exponent c') one-way functions is sufficient to construct for any c a puncturable PRF that is $2^{-\lambda^c}$-secure.

- $\mathsf{Setup}(1^\lambda)$: *Takes the unary representation of the security parameter and outputs a signing key sk and a verification key vk.*
- $\mathsf{Sign}(sk, m)$: *Takes a signing key sk and a message $m \in \mathcal{M}$, returns a signature for m.*
- $\mathsf{Ver}(vk, m, sig)$: *Takes the public verification key vk, a message $m \in \mathcal{M}$ and an alleged signature sig for m, outputs 1 if s is a valid signature for m.*

Correctness *We require the following for all messages $m \in \mathcal{M}$.*

$$\Pr\left[\mathsf{Ver}(vk, m, s) = 1 : \begin{array}{c} sk, vk \leftarrow \mathsf{Setup}(1^\lambda) \\ s \leftarrow \mathsf{Sign}(sk, m) \end{array}\right] = 1.$$

Adaptive existential-unforgeability security under chosen message attack (EUF-CMA) *We require that any QPT adversary $\mathcal{A}$ wins the following game with negligible probability.*

1. *The challenger samples the keys $sk, vk \leftarrow \mathsf{Setup}(1)$.*
2. *The adversary $\mathcal{A}$ receives vk, and interacts with the signing oracle by sending classical messages $m \in \mathcal{M}$ and receiving the corresponding signatures that is computed by the challenger as $sig \leftarrow \mathsf{Sign}(sk, m)$.*
3. *The adversary $\mathcal{A}$ outputs a message $m^* \in \mathcal{M}$ that it has not queried the oracle with and a forged signature sig^*.*
4. *The challenger outputs 1 if and only if $\mathsf{Ver}(vk, m^*, sig^*) = 1$.*

If the adversary $\mathcal{A}$ is required to output the message m^* before receiving vk, we call it selective EUF-CMA *security.*

3.4 Indistinguishability Obfuscation

In this section, we recall indistinguishability obfuscation (iO).

Definition 4 (Indistinguishability Obfuscation)**.** *An indistinguishability obfuscation (iO) scheme for a class of circuits $\mathcal{C} = \{\mathcal{C}_\lambda\}_\lambda$ is an efficient algorithm $i\mathcal{O}$ that satisfies the following.*

Correctness. *For all $\lambda \in \mathbb{N}^+, C \in \mathcal{C}_\lambda$ and all inputs x to C,*

$$\Pr\left[\tilde{C}(x) = C(x) : \tilde{C} \leftarrow i\mathcal{O}(1^\lambda, C)\right] = 1.$$

Security. *Let $\mathcal{B}$ be any QPT algorithm that outputs two circuits $C_0, C_1 \in \mathcal{C}$ of the same size, along with quantum auxiliary information $\mathsf{R}_{\mathsf{aux}}$, such that $\Pr\left[\forall x\ C_0(x) = C_1(x) : (C_0, C_1, \mathsf{R}_{\mathsf{aux}}) \leftarrow \mathcal{B}(1^\lambda)\right] \geq 1 - \mathsf{negl}(\lambda)$. Then, for any QPT adversary $\mathcal{A}$,*

$$\left|\Pr\left[\mathcal{A}(i\mathcal{O}(1^\lambda, C_0), \mathsf{R}_{\mathsf{aux}}) = 1 : (C_0, C_1, \mathsf{R}_{\mathsf{aux}}) \leftarrow \mathcal{B}(1^\lambda)\right] - \Pr\left[\mathcal{A}(i\mathcal{O}(1^\lambda, C_1), \mathsf{R}_{\mathsf{aux}}) = 1 : (C_0, C_1, \mathsf{R}_{\mathsf{aux}}) \leftarrow \mathcal{B}(1^\lambda)\right]\right| \leq \mathsf{negl}(\lambda).$$

For subexponentially secure obfuscation, we require security for circuit samplers $\mathcal{B}$ such that $\Pr\left[\forall x\ C_0(x) = C_1(x) : (C_0, C_1, \mathsf{R}_{\mathsf{aux}}) \leftarrow \mathcal{B}(1^\lambda)\right] = 1$.

Fix an arbitrary constant $c > 0$ and set $n := \lceil \lambda^{c/3} \rceil$.

3.5 Subspace-Hiding Obfuscation

Definition 5 (Subspace-hiding obfuscation)**.** *Let $N = N(\lambda)$ and $d_0 = d_0(\lambda) < d_1 = d_1(\lambda) \le N(\lambda)$ be polynomially bounded. A subspace-hiding obfuscator shO is a PPT algorithm that, on input 1^λ and a subspace $S \le \mathbb{F}_2^N$ of dimension d_0 or d_1, outputs a circuit $\widehat{S} \leftarrow \mathsf{shO}(1^\lambda, S)$.*

Correctness. *For every such subspace S, except with negligible probability over the randomness of shO, $\widehat{S}(x) = 1$ if and only if $x \in S$ for every $x \in \mathbb{F}_2^N$.*

Security. *For every QPT algorithm $\mathcal{B}$ that outputs a d_0-dimensional subspace $S_0 \le \mathbb{F}_2^N$ together with quantum auxiliary information R, the following distributions are computationally indistinguishable:*

$$\left\{(R, \widehat{S}_0) : (S_0, R) \leftarrow \mathcal{B}(1^\lambda),\ \widehat{S}_0 \leftarrow \mathsf{shO}(1^\lambda, S_0)\right\} \approx_c \left\{(R, \widehat{S}_1) : \begin{array}{l} (S_0, R) \leftarrow \mathcal{B}(1^\lambda), \\ S_1 \xleftarrow{\$} \{S \le \mathbb{F}_2^N : \dim(S) = d_1,\ S_0 \subseteq S\}, \\ \widehat{S}_1 \leftarrow \mathsf{shO}(1^\lambda, S_1) \end{array}\right\}. \tag{9}$$

Theorem 8 (Subspace-hiding obfuscation from $i\mathcal{O}$ and one-way functions [Zha21, CLLZ21, ABH+26])**.** *Assume the existence of post-quantum polynomially secure $i\mathcal{O}$ and post-quantum one-way functions. Then, for every polynomially bounded $n = n(\lambda) \in \omega(\log \lambda)$, there exists a post-quantum secure subspace-hiding obfuscator for $N = 4n$, $d_0 = 2n$, and $d_1 = 3n$. These are the parameters used in Theorem 21.*

3.6 Universal Hash Functions and Average Conditional Min-Entropy

Definition 6 (Average conditional min-entropy)**.** *For jointly distributed classical random variables X, Z, define*

$$\widetilde{H}_\infty(X \mid Z) := -\log\left(\sum_z \Pr[Z = z] \max_x \Pr[X = x \mid Z = z]\right). \tag{10}$$

When Z is empty, this is the usual min-entropy of X.

Definition 7 (Universal hash family)**.** *A family $\mathcal{H}$ of functions from $\{0,1\}^\ell$ to $\{0,1\}^m$ is universal if it is efficiently sampleable and, for every distinct $x, x' \in \{0,1\}^\ell$, $\Pr_{h \xleftarrow{\$} \mathcal{H}}[h(x) = h(x')] \le 2^{-m}$.*

Theorem 9 (Universal hash families)**.** *For every polynomially bounded $\ell = \ell(\lambda)$ and $m = m(\lambda)$, there exists an efficiently sampleable universal hash family from $\{0,1\}^\ell$ to $\{0,1\}^m$. In Theorem 23, we use such a family from $\mathcal{X}_\lambda = \{0,1\}^{\ell_{\mathsf{in}}(\lambda)}$ to $\{0,1\}^{8n}$.*

Lemma 1 (Leakage bound and leftover hash lemma)**.** *Let X, Y, Z be jointly distributed classical random variables. If $|\mathrm{Supp}(Y)| \le 2^q$, then $\widetilde{H}_\infty(X \mid Z, Y) \ge \widetilde{H}_\infty(X \mid Z) - q$. Moreover, let $\mathcal{H}$ be a universal hash family from the support of X to $\{0,1\}^m$, let $h \xleftarrow{\$} \mathcal{H}$ be independent of (X, Z), and let $U_m \xleftarrow{\$} \{0,1\}^m$ be independent of (h, X, Z). If $\widetilde{H}_\infty(X \mid Z) \ge m + 2\log(1/\varepsilon)$, then $(h, Z, h(X))$ and (h, Z, U_m) are within statistical distance at most ε.*

3.7 Injective/Lossy Functions

Definition 8 (Injective/lossy functions)**.** *An injective/lossy function family* $\mathsf{LF} = (\mathsf{LF.Gen}_{\mathsf{inj}}, \mathsf{LF.Gen}_{\mathsf{loss}})$ *with domain* $\mathcal{X}_\lambda$ *and range* $\mathcal{T}_\lambda$ *consists of two PPT algorithms that output descriptions of deterministic functions* $L : \mathcal{X}_\lambda \to \mathcal{T}_\lambda$*. We require the following.*

1. *Every* $L \leftarrow \mathsf{LF.Gen}_{\mathsf{inj}}(1^\lambda)$ *is injective.*
2. *There exists a polynomially bounded function* $\ell_{\mathsf{img}} = \ell_{\mathsf{img}}(\lambda)$ *such that every* $L \leftarrow \mathsf{LF.Gen}_{\mathsf{loss}}(1^\lambda)$ *satisfies* $|\mathrm{Im}(L)| \leq 2^{\ell_{\mathsf{img}}}$*.*
3. *The distributions* $\{\mathsf{LF.Gen}_{\mathsf{inj}}(1^\lambda)\}_\lambda$ *and* $\{\mathsf{LF.Gen}_{\mathsf{loss}}(1^\lambda)\}_\lambda$ *are computationally indistinguishable against QPT algorithms.*

Theorem 10 (Injective/lossy functions from LWE [ABH+26])**.** *Assume the quantum hardness of LWE. Then, for every constant* $c > 0$ *and every polynomially bounded input length* $\ell_{\mathsf{in}} = \ell_{\mathsf{in}}(\lambda)$ *satisfying* $\ell_{\mathsf{in}}(\lambda) \geq \lambda^c$*, there exists a post-quantum secure injective/lossy function family with domain* $\{0,1\}^{\ell_{\mathsf{in}}(\lambda)}$ *for which* $\ell_{\mathsf{img}} \leq \lambda^{2c/3}$*. This is the parameter setting used in the proof of Theorem 23.*

3.8 Unclonable Encryption

Next, we define (one-time) unclonable encryption for single bit messages.

Definition 9 (Unclonable encryption for bits)**.** *An unclonable encryption scheme for bits is a tuple* $\mathsf{UE} = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ *with the following syntax.*

- $k \leftarrow \mathsf{KeyGen}(1^\lambda)$ *outputs a classical key* $k \in \mathcal{K}_\lambda$*.*
- $\rho_{\mathsf{ct}} \leftarrow \mathsf{Enc}(k, m)$*, on input* $m \in \{0,1\}$*, outputs a quantum ciphertext in a finite-dimensional Hilbert space* $\mathcal{H}_{\mathsf{ct}}$*.*
- $m' \leftarrow \mathsf{Dec}(k, \rho_{\mathsf{ct}})$ *outputs a bit.*

For every $m \in \{0,1\}$*, correctness requires*

$$\Pr\left[m' = m : \begin{array}{c} k \leftarrow \mathsf{KeyGen}(1^\lambda), \\ \rho_{\mathsf{ct}} \leftarrow \mathsf{Enc}(k, m), \\ m' \leftarrow \mathsf{Dec}(k, \rho_{\mathsf{ct}}) \end{array}\right] \geq 1 - \mathsf{negl}(\lambda). \tag{11}$$

For arbitrary finite-dimensional Hilbert spaces $\mathcal{H}_\mathcal{B}$ and $\mathcal{H}_\mathcal{C}$, let an adversary $\mathcal{A}_{\mathsf{UE}}$ consist of a quantum channel

$$\Phi_\mathcal{A} : \mathcal{L}(\mathcal{H}_{\mathsf{ct}}) \longrightarrow \mathcal{L}(\mathcal{H}_\mathcal{B} \otimes \mathcal{H}_\mathcal{C}) \tag{12}$$

and, for every $k \in \mathcal{K}_\lambda$, binary POVMs $\{B_0^k, B_1^k\}$ on $\mathcal{H}_\mathcal{B}$ and $\{C_0^k, C_1^k\}$ on $\mathcal{H}_\mathcal{C}$. Consider the following experiment.

$\mathsf{UE\text{-}Unclonable\text{-}Indistinguishable}_{\mathsf{UE},\mathcal{A}_{\mathsf{UE}}}(1^\lambda)$:

1. Ch samples $k \leftarrow \mathsf{KeyGen}(1^\lambda)$ and $m \xleftarrow{\$} \{0,1\}$, computes $\rho_{\mathsf{ct}} \leftarrow \mathsf{Enc}(k, m)$, and sends ρ_{ct} to $\mathcal{A}$ while retaining k.
2. $\mathcal{A}$ applies $\Phi_\mathcal{A}$ and sends the resulting registers to $\mathcal{B}$ and $\mathcal{C}$.
3. After the split, Ch sends k to both $\mathcal{B}$ and $\mathcal{C}$.

4. $\mathcal{B}$ measures with $\{B_0^k, B_1^k\}$ and outputs $\widehat{m}_{\mathcal{B}}$, while $\mathcal{C}$ measures with $\{C_0^k, C_1^k\}$ and outputs $\widehat{m}_{\mathcal{C}}$.

5. The output of the experiment is 1 if $\widehat{m}_{\mathcal{B}} = \widehat{m}_{\mathcal{C}} = m$.

Definition 10 (ϵ-unclonable indistinguishable security)**.** *For $\epsilon = \epsilon(\lambda) \geq 0$, define the winning advantage of $\mathcal{A}_{UE}$ by*

$$\mathrm{Adv}^{\mathsf{UI}}_{\mathsf{UE},\mathcal{A}_{UE}}(\lambda) := \Pr\left[\mathsf{UE\text{-}UI}_{\mathsf{UE},\mathcal{A}_{UE}}(1^\lambda) = 1\right] - \frac{1}{2}. \tag{13}$$

The scheme UE satisfies ϵ-unclonable indistinguishable security if, for every (potentially unbounded) adversary $\mathcal{A}_{UE}$ as above,

$$\mathrm{Adv}^{\mathsf{UI}}_{\mathsf{UE},\mathcal{A}_{UE}}(\lambda) \leq \epsilon(\lambda), \tag{14}$$

or equivalently,

$$\Pr\left[\mathsf{UE\text{-}UI}_{\mathsf{UE},\mathcal{A}_{UE}}(1^\lambda) = 1\right] \leq \frac{1}{2} + \epsilon(\lambda). \tag{15}$$

3.9 Quantum Information

We recall some facts about quantum information that we use later.

Gentle measurement and rewinding. The almost-as-good-as-new lemma shows that a measurement that accepts with high probability can be gently rewound.

Lemma 2 (Almost As Good As New Lemma [Aar16]; see also [Win99])**.** *Let $d, d' \in \mathbb{N}^+$, let ρ be a mixed state over $\mathbb{C}^d$, and let U be a unitary over $\mathbb{C}^d \otimes \mathbb{C}^{d'}$. Let $(\Pi_0, \Pi_1 = I - \Pi_0)$ be a projective measurement over $\mathbb{C}^d \otimes \mathbb{C}^{d'}$. Consider the measurement that appends an ancillary register initialized to $|0\rangle\langle 0|$, applies U, and then measures (Π_0, Π_1). Suppose that, for some $\epsilon \in [0,1]$,*

$$\mathrm{Tr}\left[\Pi_0 U\left(\rho \otimes |0\rangle\langle 0|\right) U^\dagger\right] = 1 - \epsilon.$$

Define

$$\widetilde{\rho} := \mathrm{Tr}_{\mathbb{C}^{d'}}\left[U^\dagger\left(\sum_{b\in\{0,1\}} \Pi_b U\left(\rho \otimes |0\rangle\langle 0|\right) U^\dagger \Pi_b\right) U\right].$$

Then,

$$\|\rho - \widetilde{\rho}\|_{Tr} \leq \sqrt{\epsilon}.$$

We refer to the operation of undoing U after the measurement and then tracing out the ancillary register as *rewinding* the measured register. Thus, $\widetilde{\rho}$ in Lemma 2 is the state obtained after rewinding.

Projective and threshold implementations.

Definition 11 (Projective implementation [Zha20])**.** *Let $\mathcal{P} = (P, I - P)$ be a binary-outcome POVM. A projective implementation of $\mathcal{P}$ is a projective measurement $\mathcal{E} = (\Pi_p)_{p\in S}$, for a finite set $S \subseteq [0,1]$, such that the following measurement is equivalent to $\mathcal{P}$: apply $\mathcal{E}$ to obtain p, and output 1 with probability p and 0 otherwise. We denote the unique projective implementation of $\mathcal{P}$ by $\mathsf{ProjImp}(\mathcal{P})$.*

We use the following efficient, gapped form of approximate threshold implementation for a general efficient POVM. The fixed running-time and reference-system guarantees are recorded explicitly because both are needed in our reductions.

Lemma 3 (Fixed-time gapped approximate threshold instrument [Zha20, ALL+21])**.** *Let $\mathcal{P} = (P, I-P)$ be an efficient binary-outcome POVM on a register A, where an efficient implementation means a uniform polynomial-size unitary dilation (with its inverse) followed by a one-qubit projective measurement. Let $0 \le s < c \le 1$, where $\Delta := c - s$ is inverse polynomial, and let $\eta \in (0,1]$ be inverse polynomial. There is a uniform, fixed-polynomial-time two-outcome instrument*

$$\mathsf{GATI}_{\eta}^{c,s}(\mathcal{P}) = \big(\mathcal{A}, \mathcal{R}\big)$$

with the following properties. Write

$$\Pi_{\mathsf{hi}} := \mathbf{1}[P \ge c], \qquad \Pi_{\mathsf{lo}} := \mathbf{1}[P \le s].$$

For every subnormalized state ρ_{AR}, with an arbitrary reference register R, and $\sigma := (\mathcal{A} \otimes \mathrm{Id}_{\mathsf{R}})(\rho)$, we have

$$\mathrm{Tr}[(\Pi_{\mathsf{hi}} \otimes I)\rho] - \eta \,\mathrm{Tr}(\rho) \le \mathrm{Tr}(\sigma), \tag{16}$$

$$\mathrm{Tr}(\sigma) \le \mathrm{Tr}[(\mathbf{1}[P > s] \otimes I)\rho] + \eta \,\mathrm{Tr}(\rho), \tag{17}$$

$$\mathrm{Tr}[(\Pi_{\mathsf{lo}} \otimes I)\sigma] \le \eta \,\mathrm{Tr}(\rho). \tag{18}$$

Consequently, for

$$\overline{\sigma} := (\mathbf{1}[P > s] \otimes I)\sigma(\mathbf{1}[P > s] \otimes I),$$

we have

$$\mathrm{Tr}(\overline{\sigma}) \ge \mathrm{Tr}(\sigma) - \eta \,\mathrm{Tr}(\rho), \qquad \|\sigma - \overline{\sigma}\|_1 \le 3\sqrt{\eta}\ \mathrm{Tr}(\rho). \tag{19}$$

The running time is

$$T_{\mathcal{P}} \cdot \mathsf{poly}\big(\Delta^{-1}, \log(1/\eta), \eta^{-1}\big)\,,$$

where $T_{\mathcal{P}}$ is the running time of the dilation of $\mathcal{P}$. The same circuit and bounds apply uniformly for every reference register.

Proof. Purify all classical coins and deferred measurements in the implementation of $\mathcal{P}$. To recall why the cited threshold construction applies to a general efficient POVM, let Π_0 project the dilation workspace onto its initial state and let Π_1 be the accepting projector pulled back through the dilation unitary. On the image of Π_0,

$$\Pi_0 \Pi_1 \Pi_0 = P.$$

The standard two-projector (Jordan-block) construction underlying the approximate projective implementation of Zhandry [Zha20] and the approximate threshold implementation of Aaronson et al. [ALL+21] therefore measures the spectral parameter of P and returns the dilation workspace to its initial state. This gives an instrument on A, not merely on the dilation space.

Set $\xi := \Delta/4$ and run that approximate threshold implementation at threshold $c - \xi$, with shift and almost-projectivity errors at most $\eta/20$. Its shift guarantee compares its acceptance probability from below with the exact threshold projector $\mathbf{1}[P \ge c]$ and from above with $\mathbf{1}[P > c - 2\xi] \preceq \mathbf{1}[P > s]$, up to error $\eta/20$. Its post-measurement guarantee places all but $\eta/10$ of the accepted subnormalized state above threshold $c - 3\xi > s$. These are operator and local post-state guarantees on A. For an input entangled with R, the reduced accepted state on A is exactly $\mathcal{A}(\rho_{\mathsf{A}})$; hence the same probability and low-subspace bounds hold after tensoring the construction with I_{R}.

The cited implementation has expected running time $T_{\mathcal{P}}\mathsf{poly}(\Delta^{-1}, \log(1/\eta))$, uniformly over its input state. Fix a corresponding polynomial upper bound T_{exp} and truncate after $20T_{\mathsf{exp}}/\eta$ steps, outputting reject if the cutoff is reached. Markov's inequality bounds the cutoff probability by $\eta/20$ for every input, including one entangled with a reference. Absorbing the approximation and truncation errors proves Equations (16) to (18) and gives the claimed fixed polynomial running time. Finally, Equation (19) follows from Equation (18) and the gentle-measurement inequality, after normalizing ρ; the displayed constant also absorbs the discarded low-subspace block. □

Definition 12 (Threshold implementation [Zha20, ALL+21])**.** *Let $\mathcal{P}$ and $\mathsf{ProjImp}(\mathcal{P}) = (\Pi_p)_{p\in S}$ be as above. For a threshold $\tau \in [0,1]$, the threshold implementation of $\mathcal{P}$ with threshold τ is the binary projective measurement*

$$\mathsf{TI}_\tau(\mathcal{P}) := \left(\sum_{p\geq\tau} \Pi_p,\ I - \sum_{p\geq\tau} \Pi_p\right).$$

When clear from context, we also use $\mathsf{TI}_\tau(\mathcal{P})$ to denote its projector corresponding to outcome 1.

Part I
Definitional Work

4 Single-Decryptor Encryption Security Definitions

In this section, we present security definitions for single-decryptor encryption (SDE). We first define syntax and correctness. We then recall the prior security definitions in Section 4.1 and present our new definition in Section 4.2. Finally, in Section 4.3, we discuss the relationships between these notions.

Now we present the syntax.

Definition 13 (Single-Decryptor Encryption (SDE))**.** *A public-key single-decryptor encryption scheme* $\mathbf{SDE} = (\mathsf{GenKey}, \mathsf{Enc}, \mathsf{Dec})$ *for the message space* $\mathcal{M}$ *consists of the following efficient, possibly quantum, algorithms.*

- $\mathsf{GenKey}(1^\lambda)$*: Takes as input a security parameter and outputs a classical public key* pk *and a quantum decryption key* sk*.*
- $\mathsf{Enc}(\mathrm{pk}, \mathrm{m})$*: Takes as input a public-key* pk *and a message* $\mathrm{m} \in \mathcal{M}$ *and outputs a classical ciphertext* ct*.*
- $\mathsf{Dec}(\mathsf{sk}, \mathrm{ct})$*: Takes as input a decryption key* sk *and a ciphertext* ct*, and outputs a message* $\mathrm{m} \in \mathcal{M}$ *or* $\bot$*.*

We require the following correctness guarantee.

Correctness: *For any message* $\mathrm{m} \in \mathcal{M}$*, there exists a negligible function* $\mathsf{negl}(\lambda)$ *such that*

$$\Pr\left[\mathrm{m}' = \mathrm{m} : \begin{array}{c}(\mathrm{pk}, \mathsf{sk}) \leftarrow \mathsf{GenKey}(1^\lambda) \\ \mathrm{ct} \leftarrow \mathsf{Enc}(\mathrm{pk}, \mathrm{m}) \\ \mathrm{m}' \leftarrow \mathsf{Dec}(\mathsf{sk}, \mathrm{ct})\end{array}\right] \geq 1 - \mathsf{negl}(\lambda).$$

4.1 Prior Definitions

We recall the prior security definitions for SDE. Each game is between a challenger and an adversary tuple $\mathcal{A} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$. The adversary $\mathcal{A}$ receives the quantum decryption key and splits it into two, possibly entangled, registers for $\mathcal{B}$ and $\mathcal{C}$, who do not communicate after receiving their registers.

We begin with the two search-type notions.

Definition 14 (Independent-Challenge Search Anti-Piracy [CLLZ21, KY25])**.** *Let* $\mathbf{SDE}$ *be an SDE scheme with message space* $\mathcal{M}$*, and let* $\mathcal{A} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ *be an adversary tuple. Consider the following game.*

$\underline{\mathsf{SDE\text{-}IND\text{-}SEARCH}_{\mathbf{SDE},\mathcal{A}}(1^\lambda)}$

1. *The challenger samples* $(\mathrm{pk}, \mathsf{sk}) \leftarrow \mathbf{SDE}.\mathsf{GenKey}(1^\lambda)$ *and submits* $(\mathrm{pk}, \mathsf{sk})$ *to* $\mathcal{A}$*.*
2. $\mathcal{A}$ *outputs a pair of registers* $(\mathsf{R}_1, \mathsf{R}_2)$*. The challenger submits* R_1 *to* $\mathcal{B}$ *and* R_2 *to* $\mathcal{C}$*.*
3. *The challenger independently samples* $\mathrm{m}_\mathcal{B}, \mathrm{m}_\mathcal{C} \leftarrow \mathcal{M}$ *and, in two independent executions of encryption, computes* $\mathrm{ct}_\mathcal{B} \leftarrow \mathbf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m}_\mathcal{B})$ *and* $\mathrm{ct}_\mathcal{C} \leftarrow \mathbf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m}_\mathcal{C})$*.*

4. *The challenger submits* $\mathrm{ct}_{\mathcal{B}}$ *to* $\mathcal{B}$ *and* $\mathrm{ct}_{\mathcal{C}}$ *to* $\mathcal{C}$*. Then,* $\mathcal{B}$ *outputs* $\widetilde{\mathrm{m}}_{\mathcal{B}}$ *and* $\mathcal{C}$ *outputs* $\widetilde{\mathrm{m}}_{\mathcal{C}}$*.*

5. *The challenger outputs* 1 *if and only if* $\widetilde{\mathrm{m}}_{\mathcal{B}} = \mathrm{m}_{\mathcal{B}}$ *and* $\widetilde{\mathrm{m}}_{\mathcal{C}} = \mathrm{m}_{\mathcal{C}}$*.*

SDE *satisfies* independent-challenge search anti-piracy security *if, for every QPT adversary* $\mathcal{A}$*,*

$$\Pr\left[\mathsf{SDE\text{-}IND\text{-}SEARCH}_{\mathbf{SDE},\mathcal{A}}(1^\lambda) = 1\right] \leq \frac{1}{|\mathcal{M}|} + \mathsf{negl}(\lambda).$$

Definition 15 (Identical-Challenge Search Anti-Piracy [KY25])**.** *Let* **SDE** *be an SDE scheme with message space* $\mathcal{M}$*, and let* $\mathcal{A} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ *be an adversary tuple. Consider the following game.*

$\mathsf{SDE\text{-}IDEN\text{-}SEARCH}_{\mathbf{SDE},\mathcal{A}}(1^\lambda)$

1. *The challenger samples* $(\mathrm{pk}, s\mathcal{k}) \leftarrow \mathbf{SDE}.\mathsf{GenKey}(1^\lambda)$ *and submits* $(\mathrm{pk}, s\mathcal{k})$ *to* $\mathcal{A}$*.*

2. $\mathcal{A}$ *outputs a pair of registers* $(\mathsf{R}_1, \mathsf{R}_2)$*. The challenger submits* R_1 *to* $\mathcal{B}$ *and* R_2 *to* $\mathcal{C}$*.*

3. *The challenger samples* $\mathrm{m} \leftarrow \mathcal{M}$ *and computes a single ciphertext* $\mathrm{ct} \leftarrow \mathbf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m})$*.*

4. *The challenger submits this same ciphertext* ct *to both* $\mathcal{B}$ *and* $\mathcal{C}$*. Then,* $\mathcal{B}$ *outputs* $\widetilde{\mathrm{m}}_{\mathcal{B}}$ *and* $\mathcal{C}$ *outputs* $\widetilde{\mathrm{m}}_{\mathcal{C}}$*.*

5. *The challenger outputs* 1 *if and only if* $\widetilde{\mathrm{m}}_{\mathcal{B}} = \mathrm{m}$ *and* $\widetilde{\mathrm{m}}_{\mathcal{C}} = \mathrm{m}$*.*

SDE *satisfies* identical-challenge search anti-piracy security *if, for every QPT adversary* $\mathcal{A}$*,*

$$\Pr\left[\mathsf{SDE\text{-}IDEN\text{-}SEARCH}_{\mathbf{SDE},\mathcal{A}}(1^\lambda) = 1\right] \leq \frac{1}{|\mathcal{M}|} + \mathsf{negl}(\lambda).$$

We next recall the two CPA-type notions. In the independent-challenge game, the adversary may choose a different message pair for each party ([KY25]). The original definition of Coladangelo et al. ([CLLZ21]) requires the two pairs to be equal.

Definition 16 (Independent-Challenge CPA Anti-Piracy [CLLZ21, KY25])**.** *Let* **SDE** *be an SDE scheme with message space* $\mathcal{M}$*, and let* $\mathcal{A} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ *be an adversary tuple. Consider the following game.*

$\mathsf{SDE\text{-}IND\text{-}CPA}_{\mathbf{SDE},\mathcal{A}}(1^\lambda)$

1. *The challenger samples* $(\mathrm{pk}, s\mathcal{k}) \leftarrow \mathbf{SDE}.\mathsf{GenKey}(1^\lambda)$ *and submits* $(\mathrm{pk}, s\mathcal{k})$ *to* $\mathcal{A}$*.*

2. $\mathcal{A}$ *outputs message pairs* $(\mathrm{m}_{\mathcal{B},0}, \mathrm{m}_{\mathcal{B},1}) \in \mathcal{M}^2$ *and* $(\mathrm{m}_{\mathcal{C},0}, \mathrm{m}_{\mathcal{C},1}) \in \mathcal{M}^2$*, and a pair of registers* $(\mathsf{R}_1, \mathsf{R}_2)$*. The challenger submits* R_1 *to* $\mathcal{B}$ *and* R_2 *to* $\mathcal{C}$*.*

3. *The challenger samples* $b_{\mathcal{B}}, b_{\mathcal{C}} \leftarrow \{0,1\}$ *independently. In two independent executions of encryption, it computes* $\mathrm{ct}_{\mathcal{B}} \leftarrow \mathbf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m}_{\mathcal{B},b_{\mathcal{B}}})$ *and* $\mathrm{ct}_{\mathcal{C}} \leftarrow \mathbf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m}_{\mathcal{C},b_{\mathcal{C}}})$*.*

4. *The challenger submits* $\mathrm{ct}_{\mathcal{B}}$ *to* $\mathcal{B}$ *and* $\mathrm{ct}_{\mathcal{C}}$ *to* $\mathcal{C}$*. Then,* $\mathcal{B}$ *outputs* $\widetilde{b}_{\mathcal{B}}$ *and* $\mathcal{C}$ *outputs* $\widetilde{b}_{\mathcal{C}}$*.*

5. *The challenger outputs* 1 *if and only if* $\widetilde{b}_{\mathcal{B}} = b_{\mathcal{B}}$ *and* $\widetilde{b}_{\mathcal{C}} = b_{\mathcal{C}}$*.*

SDE *satisfies* independent-challenge CPA anti-piracy security *if, for every QPT adversary* $\mathcal{A}$*,*

$$\Pr\left[\mathsf{SDE\text{-}IND\text{-}CPA}_{\mathbf{SDE},\mathcal{A}}(1^\lambda) = 1\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

Definition 17 (Identical-Challenge CPA Anti-Piracy [GZ20])**.** *Let* **SDE** *be an SDE scheme with message space* $\mathcal{M}$*, and let* $\mathcal{A} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ *be an adversary tuple. Consider the following game.*

$\underline{\mathsf{SDE\text{-}IDEN\text{-}CPA}_{\mathbf{SDE},\mathcal{A}}(1^\lambda)}$

1. *The challenger samples* $(\mathrm{pk}, s\mathcal{k}) \leftarrow \mathbf{SDE}.\mathsf{GenKey}(1^\lambda)$ *and submits* $(\mathrm{pk}, s\mathcal{k})$ *to* $\mathcal{A}$.
2. $\mathcal{A}$ *outputs a message pair* $(\mathrm{m}_0, \mathrm{m}_1) \in \mathcal{M}^2$ *and a pair of registers* $(\mathsf{R}_1, \mathsf{R}_2)$. *The challenger submits* R_1 *to* $\mathcal{B}$ *and* R_2 *to* $\mathcal{C}$.
3. *The challenger samples* $b \leftarrow \{0,1\}$ *and computes a single ciphertext* $\mathrm{ct} \leftarrow \mathbf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m}_b)$.
4. *The challenger submits this same ciphertext* ct *to both* $\mathcal{B}$ *and* $\mathcal{C}$. *Then,* $\mathcal{B}$ *outputs* $\widetilde{b}_\mathcal{B}$ *and* $\mathcal{C}$ *outputs* $\widetilde{b}_\mathcal{C}$.
5. *The challenger outputs* 1 *if and only if* $\widetilde{b}_\mathcal{B} = b$ *and* $\widetilde{b}_\mathcal{C} = b$.

SDE *satisfies* identical-challenge CPA anti-piracy security *if, for every QPT adversary* $\mathcal{A}$,

$$\Pr\left[\mathsf{SDE\text{-}IDEN\text{-}CPA}_{\mathbf{SDE},\mathcal{A}}(1^\lambda) = 1\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

CPA$^+$ anti-piracy uses the independent-challenge CPA experiment with a stronger winning condition.

Definition 18 (CPA$^+$ Anti-Piracy Security [KY25])**.** *Let* **SDE** *be an SDE scheme and let* $\mathcal{A} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ *be an adversary tuple. The game* $\mathsf{SDE\text{-}CPA}^+_{\mathbf{SDE},\mathcal{A}}(1^\lambda)$ *is identical to the experiment in Definition 16, except that the challenger outputs* 1 *if and only if*

$$\widetilde{b}_\mathcal{B} \oplus \widetilde{b}_\mathcal{C} = b_\mathcal{B} \oplus b_\mathcal{C}.$$

SDE *satisfies* CPA$^+$ anti-piracy security *if, for every QPT adversary* $\mathcal{A}$,

$$\Pr\left[\mathsf{SDE\text{-}CPA}^+_{\mathbf{SDE},\mathcal{A}}(1^\lambda) = 1\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

We finally recall strong anti-piracy security ([CLLZ21, KY25]). A *quantum decryptor* is a pair $\mathsf{D} = (\rho, U)$, where ρ is a quantum state and U is a unitary circuit that takes a ciphertext and ρ as input and produces a classical output. We define the two tests used in the strong definitions.

CPA test. For a public key pk, a message pair $(\mathrm{m}_0, \mathrm{m}_1)$, and a quantum decryptor D, let $\mathcal{P}^{\mathsf{cpa}}_\mathsf{D}[\mathrm{pk}; \mathrm{m}_0, \mathrm{m}_1]$ denote the binary-outcome POVM corresponding to the following test: sample $b \leftarrow \{0,1\}$, compute $\mathrm{ct} \leftarrow \mathbf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m}_b)$, run $\mathsf{D}(\mathrm{ct})$, and accept if and only if the output is b. The γ-good CPA test applies

$$\mathsf{TI}_{1/2+\gamma}(\mathcal{P}^{\mathsf{cpa}}_\mathsf{D}[\mathrm{pk}; \mathrm{m}_0, \mathrm{m}_1])$$

to the quantum state of D, where for any ϵ and binary outcome POVM P, $\mathsf{TI}_\epsilon(P)$ denotes the threshold implementation (with threshold ϵ) of P, as defined in Section 3.9.

Search test. Fix a public key pk and a quantum decryptor D. Let $\mathcal{P}^{\mathsf{search}}_\mathsf{D}[\mathrm{pk}]$ denote the binary-outcome POVM corresponding to the following test. Sample $\mathrm{m} \leftarrow \mathcal{M}$, compute $\mathrm{ct} \leftarrow \mathbf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m})$, and run $\mathsf{D}(\mathrm{ct})$. Accept if and only if the output is m. The γ-good search test applies

$$\mathsf{TI}_{1/|\mathcal{M}|+\gamma}(\mathcal{P}^{\mathsf{search}}_\mathsf{D}[\mathrm{pk}])$$

to the quantum state of D.

Definition 19 (Strong CPA Anti-Piracy Security [CLLZ21, KY25])**.** *Let* $\gamma : \mathbb{N} \to [0,1]$, *let* **SDE** *be an SDE scheme with message space* $\mathcal{M}$, *and let* $\mathcal{A}$ *be an adversary. Consider the following game.*

$\underline{\mathsf{SDE\text{-}STRONG\text{-}CPA}_{\mathbf{SDE},\mathcal{A}}(1^\lambda,\gamma)}$

1. *The challenger samples* $(\mathrm{pk},s\mathcal{k}) \leftarrow \mathbf{SDE}.\mathsf{GenKey}(1^\lambda)$ *and submits* $(\mathrm{pk},s\mathcal{k})$ *to* $\mathcal{A}$.
2. $\mathcal{A}$ *outputs message pairs* $(\mathrm{m}_{\mathcal{B},0},\mathrm{m}_{\mathcal{B},1}) \in \mathcal{M}^2$ *and* $(\mathrm{m}_{\mathcal{C},0},\mathrm{m}_{\mathcal{C},1}) \in \mathcal{M}^2$*, and two possibly entangled quantum decryptors* $\mathsf{D}_\mathcal{B}$ *and* $\mathsf{D}_\mathcal{C}$.
3. *The challenger applies the* γ*-good CPA test with respect to* $(\mathrm{pk},\mathrm{m}_{X,0},\mathrm{m}_{X,1})$ *to* D_X*, for each* $X \in \{\mathcal{B},\mathcal{C}\}$*. The challenger outputs* 1 *if and only if both tests pass.*

SDE *satisfies* strong CPA anti-piracy security *if, for every inverse polynomial* γ *and every QPT adversary* $\mathcal{A}$*,*

$$\Pr\left[\mathsf{SDE\text{-}STRONG\text{-}CPA}_{\mathbf{SDE},\mathcal{A}}(1^\lambda,\gamma(\lambda))=1\right] \leq \mathsf{negl}(\lambda).$$

Definition 20 (Strong Search Anti-Piracy Security [CLLZ21])**.** *Let* $\gamma : \mathbb{N} \to [0,1]$*, let* **SDE** *be an SDE scheme with message space* $\mathcal{M}$*, and let* $\mathcal{A}$ *be an adversary. Consider the following game.*

$\underline{\mathsf{SDE\text{-}STRONG\text{-}SEARCH}_{\mathbf{SDE},\mathcal{A}}(1^\lambda,\gamma)}$

1. *The challenger samples* $(\mathrm{pk},s\mathcal{k}) \leftarrow \mathbf{SDE}.\mathsf{GenKey}(1^\lambda)$ *and submits* $(\mathrm{pk},s\mathcal{k})$ *to* $\mathcal{A}$.
2. $\mathcal{A}$ *outputs two possibly entangled quantum decryptors* $\mathsf{D}_\mathcal{B}$ *and* $\mathsf{D}_\mathcal{C}$.
3. *The challenger applies the* γ*-good search test with respect to* pk *to each decryptor. The challenger outputs* 1 *if and only if both tests pass.*

SDE *satisfies* strong search anti-piracy security *if, for every inverse polynomial* γ *and every QPT adversary* $\mathcal{A}$*,*

$$\Pr\left[\mathsf{SDE\text{-}STRONG\text{-}SEARCH}_{\mathbf{SDE},\mathcal{A}}(1^\lambda,\gamma(\lambda))=1\right] \leq \mathsf{negl}(\lambda).$$

4.2 Correlated-Challenge Decision Security

We generalize the prior definition of SDE in two ways. First we allow the challenge bits for $\mathcal{B}$ and $\mathcal{C}$ to be arbitrarily correlated as long as their marginal distributions are uniform. Secondly, we follow a Left-or-Right version of CPA game where for each freeloader $X \in \{\mathcal{B},\mathcal{C}\}$, there are polynomially many pairs (as opposed to a single pair) of challenge messages that the adversary $\mathcal{A}$ can specify, and then the challenger provides either the encryptions of the left half of the challenge message pairs or the encryptions of the right half of the challenge message pairs, depending on the challenge bit for X. Next we define an admissible sampler that specifies this allowed correlation between the challenge bits.

Definition 21 (Admissible Challenge Sampler)**.** *Fix an SDE scheme* **SDE***. An admissible challenge sampler* Samp *is an efficient, possibly quantum, algorithm with classical output such that, for every public key* pk *that* $\mathbf{SDE}.\mathsf{GenKey}(1^\lambda)$ *can output: on input* $(\mathrm{pk},\mathsf{IND})$ *it returns* $(b_\mathcal{B},b_\mathcal{C})$ *with uniform marginals, and on input* $(\mathrm{pk},\mathsf{IDEN})$ *it returns one uniform bit* b*. Equivalently, for every* $X \in \{\mathcal{B},\mathcal{C}\}$ *and* $\beta \in \{0,1\}$*,*

$$\Pr_{(b_\mathcal{B},b_\mathcal{C})\leftarrow\mathsf{Samp}(\mathrm{pk},\mathsf{IND})}[b_X=\beta]=\frac{1}{2}, \qquad \Pr_{b\leftarrow\mathsf{Samp}(\mathrm{pk},\mathsf{IDEN})}[b=\beta]=\frac{1}{2}. \tag{20}$$

The two bits returned in separate mode may otherwise be arbitrarily correlated. Identical mode has one common bit by syntax.

Definition 22 (SDE-CORR Security)**.** *Let* **SDE** *be an SDE scheme,* Samp *an admissible challenge sampler, and* $\mathcal{A} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ *a QPT adversary tuple. Consider the following game, in which one selector (or selector pair) is used throughout the batch but the message pairs may vary by coordinate.*

$\underline{\mathsf{SDE\text{-}CORR}_{\mathbf{SDE},\mathcal{A},\mathsf{Samp}}(1^\lambda)}$

1. *The challenger samples* $(\mathrm{pk}, s\mathcal{k}) \leftarrow \mathbf{SDE}.\mathsf{GenKey}(1^\lambda)$ *and submits* $(\mathrm{pk}, s\mathcal{k})$ *to* $\mathcal{A}$*.*
2. $\mathcal{A}$ *outputs* 1^N*, for* $N \geq 1$*, a mode* $\mathsf{mode} \in \{\mathsf{IND}, \mathsf{IDEN}\}$*, registers* $(\mathsf{R}_1, \mathsf{R}_2)$*, and challenge-message pairs:*
   - *if* $\mathsf{mode} = \mathsf{IND}$*, it outputs* $4N$ *messages*
   $$\left(\mathrm{m}_{\mathcal{B},i,0}, \mathrm{m}_{\mathcal{B},i,1}, \mathrm{m}_{\mathcal{C},i,0}, \mathrm{m}_{\mathcal{C},i,1}\right)_{i\in[N]} \in (\mathcal{M}^4)^N;$$
   - *if* $\mathsf{mode} = \mathsf{IDEN}$*, it outputs* $2N$ *messages*
   $$\left(\mathrm{m}_{i,0}, \mathrm{m}_{i,1}\right)_{i\in[N]} \in (\mathcal{M}^2)^N.$$
3. *The challenger sends* R_1 *to* $\mathcal{B}$ *and* R_2 *to* $\mathcal{C}$*, and creates the challenge batches as follows.*
   - *If* $\mathsf{mode} = \mathsf{IND}$*, sample* $(b_\mathcal{B}, b_\mathcal{C}) \leftarrow \mathsf{Samp}(\mathrm{pk}, \mathsf{IND})$*. For every* $X \in \{\mathcal{B}, \mathcal{C}\}$ *and* $i \in [N]$*, independently compute*
   $$\mathrm{ct}_{X,i} \leftarrow \mathbf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m}_{X,i,b_X}).$$
   - *If* $\mathsf{mode} = \mathsf{IDEN}$*, sample* $b \leftarrow \mathsf{Samp}(\mathrm{pk}, \mathsf{IDEN})$ *and set* $b_\mathcal{B} = b_\mathcal{C} := b$*. For every* $i \in [N]$*, independently compute one ciphertext*
   $$\mathrm{ct}_i \leftarrow \mathbf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m}_{i,b}),$$
   and give that same ciphertext to both parties by setting $\mathrm{ct}_{\mathcal{B},i} = \mathrm{ct}_{\mathcal{C},i} := \mathrm{ct}_i$*.*
4. *The challenger sends* $\overline{\mathrm{ct}}_X := (\mathrm{ct}_{X,1}, \ldots, \mathrm{ct}_{X,N})$ *to each* $X \in \{\mathcal{B}, \mathcal{C}\}$*.*
5. *After receiving their batches,* $\mathcal{B}$ *and* $\mathcal{C}$ *output* $\widetilde{b}_\mathcal{B}$ *and* $\widetilde{b}_\mathcal{C}$*, respectively. The challenger outputs* 1 *if and only if* $\widetilde{b}_\mathcal{B} = b_\mathcal{B}$ *and* $\widetilde{b}_\mathcal{C} = b_\mathcal{C}$*.*

We demand the output of $\mathcal{A}$ *in Item 2 to be unary, in order to make* N *polynomially bounded. Clearly,* $N = 1$ *is the usual CPA security with a single-ciphertext. The scheme* **SDE** *is* SDE-CORR secure *if, for every admissible* Samp *and every QPT* $\mathcal{A}$*,*

$$\Pr\left[\mathsf{SDE\text{-}CORR}_{\mathbf{SDE},\mathcal{A},\mathsf{Samp}}(1^\lambda) = 1\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

4.3 Hierarchy of Definitions

All statements below concerning search security assume bit-string messages of length $m(\lambda) = \omega(\log \lambda)$. Equivalently, $1/|\mathcal{M}|$ is negligible, so uniform guessing succeeds with only negligible probability in the search games.

Figure 1 summarizes the implications and separations proved or recalled below. For notions X and Y, the notation $\mathsf{X} \not\Rightarrow \mathsf{Y}$ means that some SDE scheme satisfies X but not Y.

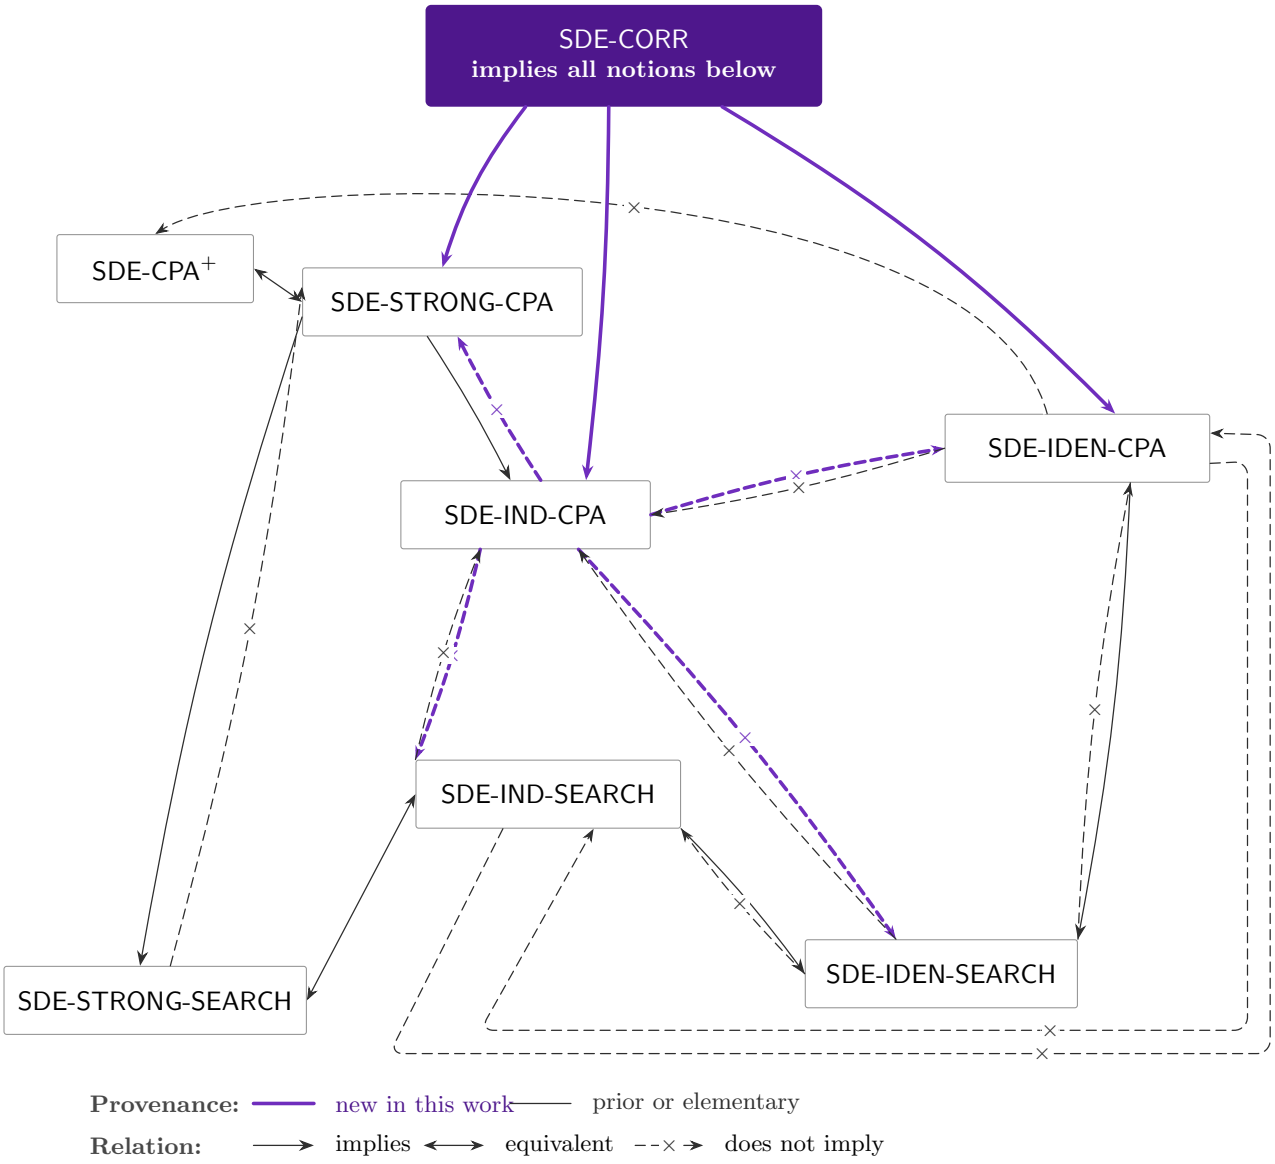

Figure 1: Hierarchy of SDE security notions under the standing assumptions of this subsection. Solid arrows denote implication, solid double-headed lines denote equivalence, and dashed directed arrows marked by a midpoint $\times$ denote failure of the implication in the arrow direction. Violet heavy lines are new in this work; black lines are prior or elementary. Unmarked directions are not claimed; in particular, the implication from SDE-CPA$^+$ to SDE-IDEN-CPA remains open.

We begin with the only implication that uses more than one challenge-message pair in SDE-CORR security (Definition 22). The proof estimates a quantum test from independently generated ciphertexts and then invokes the approximate threshold measurement of Lemma 3.

Preliminaries We first prove a perturbation bound that we need for the results below. Let A, B be Hermitian, let $\kappa > \ell$, and set $\Pi := \mathbf{1}[A \geq \kappa]$ and $Q := \mathbf{1}[B \leq \ell]$. Then

$$\|Q\Pi\| \leq \frac{\|A - B\|}{\kappa - \ell}. \tag{21}$$

Here B_Q and A_Π denote the restrictions of B and A to the ranges of Q and Π, respectively. They satisfy $B_Q(Q\Pi) - (Q\Pi)A_\Pi = -Q(A - B)\Pi$, and their spectra are separated by $\kappa - \ell$, which proves the bound. It follows that every positive operator ρ such that $\mathrm{Tr}(\rho) \leq 1$ and $\rho = \Pi\rho\Pi$ satisfies $\mathrm{Tr}(Q\rho) \leq (\|A - B\|/(\kappa - \ell))^2 \mathrm{Tr}(\rho)$. The same conclusion holds when ρ is entangled with an additional register.

We first record two immediate consequences of the definition.

Lemma 4. SDE-CORR *security implies* SDE-IND-CPA *and* SDE-IDEN-CPA *security.*

Proof. Set $N = 1$ in the SDE-CORR game of Definition 22. In separate mode, choosing independent uniform bits $b_{\mathcal{B}}, b_{\mathcal{C}}$ gives exactly the independent-challenge CPA experiment of Definition 16. In identical mode, the admissible distribution returns one uniform bit and the challenger gives a copy of the resulting ciphertext to each party, which is exactly the identical-challenge CPA experiment of Definition 17. □

We also recall the implications from prior works.

Theorem 11 (Prior relations [KY25, AKL23])**.** *The following relations hold.*

1. SDE-CPA$^+$ *and* SDE-STRONG-CPA *security are equivalent.*
2. SDE-CPA$^+$ *security implies* SDE-IND-CPA *and* SDE-IND-SEARCH *security.*
3. SDE-STRONG-SEARCH *and* SDE-IND-SEARCH *security are equivalent.*
4. SDE-IND-SEARCH *security implies* SDE-IDEN-SEARCH *security.*

Proof. The first three statements follow from [KY25, Theorem 6.24, Corollary 6.32, Remark 6.19]. The last is the independent-to-identical search implication from [KY25, Section 6.1], which uses the product-to-correlated comparison of Ananth, Kaleoglu, and Liu [AKL23]. □

The above-mentioned relations show that strong CPA implies strong search. For superlogarithmic message length, this implication is strict, in the sense that not every strong search secure SDE scheme is strong CPA secure.

Theorem 12 (Strong CPA strictly implies strong search)**.** *For superlogarithmic message length, strong-CPA security implies strong-search security. Assuming post-quantum-secure indistinguishability obfuscation (iO; Definition 4) and post-quantum-secure one-way functions, the converse does not hold. More generally, every strong-search-secure SDE scheme over a message space $\mathcal{M}_0$ satisfying $1/|\mathcal{M}_0| = \mathsf{negl}(\lambda)$ yields a strong-search-secure scheme that is not strong-CPA secure.*

Proof. By Theorem 11, strong CPA is equivalent to CPA$^+$, CPA$^+$ implies IND-SEARCH, and IND-SEARCH is equivalent to strong search.

For strictness, let Π_0 be strong-search secure over $\mathcal{M}_0$. Define Π_{clr} over $\{0,1\} \times \mathcal{M}_0$ using the same key generation. Its encryption of (a, m) is $(a, \Pi_0.\mathsf{Enc}(\mathrm{pk}, m))$, and decryption of (a, ct) returns $(a, \Pi_0.\mathsf{Dec}(sk, \mathrm{ct}))$.

Fix a decryptor D' for Π_{clr}. Construct a decryptor D for Π_0 as follows. On input ct, it prepares a uniform bit a, runs $\mathsf{D}'(a, \mathrm{ct})$ reversibly, and measures only the final output. It returns m' if the result is (a, m'), and returns a fixed symbol $\perp \notin \mathcal{M}_0$ otherwise. The two decryptors have exactly the same search-success operator. Let $K := |\mathcal{M}_0|$. Because $1/K$ is negligible, for every inverse-polynomial γ and all sufficiently large λ, we have $1/(2K) + \gamma \geq 1/K + \gamma/2$. Hence, every state that reaches the strong-search threshold for Π_{clr} with gap γ also reaches the threshold for Π_0 with gap $\gamma/2$; see Definition 20. Applying this conversion separately to the two decryptors, including when their states are entangled, transforms any strong-search attack on Π_{clr} into one on Π_0 without decreasing the probability that both threshold tests accept.

To violate strong CPA, sample $m^\star \xleftarrow{\$} \mathcal{M}_0$ and give each receiver the pair $((0, m^\star), (1, m^\star))$. Each decryptor reads the first ciphertext component and determines the challenge bit perfectly without using the quantum key. Thus, both CPA acceptance operators are the identity, and both threshold tests accept with probability one, for example when $\gamma(\lambda) = 1/\lambda$ and λ is sufficiently large. Therefore, Π_{clr} is not strong-CPA secure. Under post-quantum-secure iO and one-way functions, the SDE-CORR-secure scheme from Theorem 22 supplies the required base scheme. □

Next we show that our notion of CORR-SDE security notion implies strong CPA security.

Theorem 13 (Correlated challenges imply strong CPA)**.** *Every* SDE-CORR*-secure scheme satisfies* SDE-STRONG-CPA *security.*

Proof. We prove the contrapositive. Suppose that there are inverse polynomials γ and ε and a QPT algorithm $\mathcal{A}$, which outputs the two decryptors in the strong-CPA experiment of Definition 19, such that the experiment accepts with probability at least ε for infinitely many security parameters. We construct an adversary for the SDE-CORR game of Definition 22 with non-negligible advantage on this infinite set. Set

$$\eta := \left(\frac{\varepsilon\gamma^2}{10^6}\right)^2. \tag{22}$$

Because η is inverse polynomial, every invocation of Lemma 3 below runs in polynomial time. The constructed adversary runs $\mathcal{A}$ internally and controls all subsequent operations by its classical output; we suppress the corresponding classical registers.

Local tests. For a binary quantum test, its acceptance operator is the POVM element M for which the acceptance probability on state ρ is $\mathrm{Tr}(M\rho)$. For $X \in \{\mathcal{B}, \mathcal{C}\}$, let D_X be the decryptor produced by $\mathcal{A}$. Given a ciphertext ct, let $D_{X,y}(\mathrm{ct})$ be the measurement operator corresponding to the exact output $y \in \{0,1\}$, and define $Z_X(\mathrm{ct}) := D_{X,0}(\mathrm{ct}) - D_{X,1}(\mathrm{ct})$. For $j \in \{0,1\}$, let $Z_{X,j} := \mathbb{E}_{\mathrm{ct} \leftarrow \mathbf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m}_{X,j})}[Z_X(\mathrm{ct})]$, and set $H_X := (Z_{X,0} - Z_{X,1})/4$.

Let P_X be the acceptance operator of the CPA test in Definition 19. Let $D_{X,j,y}$ be the average of $D_{X,y}(\mathrm{ct})$ over $\mathrm{ct} \leftarrow \mathbf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m}_{X,j})$, and set $R_{X,j} := I - D_{X,j,0} - D_{X,j,1}$. Thus, $R_{X,j}$ is the average measurement operator for outputs outside $\{0,1\}$. Direct calculation gives

$$H_X - \left(P_X - \frac{I}{2}\right) = \frac{R_{X,0} + R_{X,1}}{4} \succeq 0. \tag{23}$$

Moreover, $\|H_X\| \le 1/2$, so $F_X := I/2 + H_X$ is a valid acceptance operator. It corresponds to the following binary version of the CPA test: choose $j \xleftarrow{\$} \{0,1\}$, encrypt $\mathrm{m}_{X,j}$, and run D_X; accept if the output is j, reject if it is $1-j$, and accept with probability $1/2$ on every other output. For $j=0$ and $j=1$, the acceptance operators are $I/2 + Z_{X,0}/2$ and $I/2 - Z_{X,1}/2$, respectively, and their average is F_X.

Repeating the challenge pair. In the CORR attack, $\mathcal{A}$ uses the same strong-CPA message pair in every coordinate: $\mathrm{m}_{X,i,j} := \mathrm{m}_{X,j}$ for every $i \in [N]$ and $j \in \{0,1\}$. If the common challenge bit is b, party X receives $\overline{\mathrm{ct}}_X = (\mathrm{ct}_{X,1}, \ldots, \mathrm{ct}_{X,N})$, where each $\mathrm{ct}_{X,i}$ encrypts $\mathrm{m}_{X,b}$. Define

$$\widehat{H}_{X,b} := \frac{1}{2}\left(\frac{1}{N}\sum_{i=1}^{N} Z_X(\mathrm{ct}_{X,i}) - \frac{Z_{X,0}+Z_{X,1}}{2}\right). \tag{24}$$

The second term is implemented by an independent execution of public-key encryption, and $\mathbb{E}[\widehat{H}_{X,b}] = (-1)^b H_X$. Define the following acceptance operator for a test on the N received ciphertexts:

$$G_X(\overline{\mathrm{ct}}_X) := \frac{I}{2} + \frac{\widehat{H}_{X,b}}{2}. \tag{25}$$

Although b appears in the notation because it determines the distribution of the received ciphertexts, the test does not use b. To implement it, flip a fair coin. On heads, choose $i \xleftarrow{\$} [N]$, run D_X on $\mathrm{ct}_{X,i}$, accept on output 0, reject on output 1, and accept with probability $1/2$ on every other output. The acceptance operator in this case is $(I + S_{X,b})/2$, where $S_{X,b} := N^{-1}\sum_i Z_X(\mathrm{ct}_{X,i})$. On tails, choose $j \xleftarrow{\$} \{0,1\}$, independently encrypt $\mathrm{m}_{X,j}$, run D_X, and reverse the acceptance rule. The average acceptance operator in this case is $(I - \overline{Z}_X)/2$, where $\overline{Z}_X := (Z_{X,0} + Z_{X,1})/2$. Averaging the two cases gives $I/2 + (S_{X,b} - \overline{Z}_X)/4 = G_X(\overline{\mathrm{ct}}_X)$. Thus, $G_X(\overline{\mathrm{ct}}_X)$ defines an efficient binary test, and $\mathbb{E}[G_X(\overline{\mathrm{ct}}_X)] = I/2 + (-1)^b H_X/2$.

We may assume that $\gamma \le 1/2$, because otherwise the strong-CPA threshold exceeds one and the premise is vacuous. Let q_X be the number of qubits in the state and initialized workspace of D_X. Applying Tropp's matrix concentration bound to the independent Hermitian operators above gives [Tro12]

$$\Pr\left[\|\widehat{H}_{X,b} - (-1)^b H_X\| > \frac{\gamma}{100}\right] \le 2^{q_X+1}\exp\left(-\frac{N\gamma^2}{25000}\right). \tag{26}$$

Choose the following common power-of-two number of challenge pairs:

$$N := 2^{\lceil \log_2(10^5(\max_X q_X + \lambda)/\gamma^2)\rceil}. \tag{27}$$

The probability that Equation (26) fails for either party is then negligible.

The remainder of the proof applies the approximate threshold measurement of Lemma 3, with error η, to the efficient tests with acceptance operators P_X, F_X, and $G_X(\overline{\mathrm{ct}}_X)$. Each test has a reversible implementation using fresh randomness, the public encryption circuit, and the decryptor circuit.

First threshold measurement. Let ρ be the joint decryptor state, define $\Pi_X^+ := \mathbf{1}[P_X \ge 1/2 + \gamma]$, and set $p_{\mathrm{str}} := \mathrm{Tr}\left[(\Pi_{\mathcal{B}}^+ \otimes \Pi_{\mathcal{C}}^+)\rho\right]$. Averaging over the classical output of $\mathcal{A}$, we have $\mathbb{E}[p_{\mathrm{str}}] \ge \varepsilon$.

For each party, apply the threshold measurement of Lemma 3 to P_X, with completeness threshold $1/2+\gamma$ and soundness threshold $1/2+3\gamma/4$. Let $\mathcal{A}_X^P$ be the map corresponding to acceptance,

let $\sigma_P := (\mathcal{A}^P_{\mathcal{B}} \otimes \mathcal{A}^P_{\mathcal{C}})(\rho)$, and let $p_P := \mathrm{Tr}(\sigma_P)$. Completeness remains valid when the measured register is entangled with the other party; applying it successively to the two parties gives

$$p_P \geq p_{\mathrm{str}} - 2\eta. \tag{28}$$

Set $\Pi_X := \mathbf{1}[P_X > 1/2 + 3\gamma/4]$, $\Pi := \Pi_{\mathcal{B}} \otimes \Pi_{\mathcal{C}}$, and $\overline{\sigma}_P := \Pi \sigma_P \Pi$. The post-measurement guarantee of Lemma 3 places at most η probability outside each local projector. Therefore, $\mathrm{Tr}(\overline{\sigma}_P) \geq p_P - 2\eta$, and the gentle measurement bound from Section 3.9 gives $\|\sigma_P - \overline{\sigma}_P\|_1 \leq 3\sqrt{2\eta}$.

Second threshold measurement. We now pass from states supported where $P_X > 1/2 + 3\gamma/4$ to states with non-negligible support where $H_X \geq \gamma/2$. Let $\Omega_X := \mathbf{1}[H_X \geq \gamma/2]$. By Equation (23), $\Pi_X H_X \Pi_X \succeq (3\gamma/4)\Pi_X$. Since $\|H_X\| \leq 1/2$, we also have $\Pi_X H_X \Pi_X \preceq (\gamma/2)\Pi_X + (1/2 - \gamma/2)\Pi_X \Omega_X \Pi_X$. Setting $c := \gamma/(2(1-\gamma))$, which is at least $\gamma/2$, gives

$$\Pi_X \Omega_X \Pi_X \succeq c \Pi_X. \tag{29}$$

Apply the same threshold measurement to $F_X = I/2 + H_X$, now with completeness threshold $1/2 + \gamma/2$ and soundness threshold $1/2 + \gamma/4$. Let $\mathcal{A}^H_X$ be its acceptance map, let $\sigma_H := (\mathcal{A}^H_{\mathcal{B}} \otimes \mathcal{A}^H_{\mathcal{C}})(\sigma_P)$, and let $p_H := \mathrm{Tr}(\sigma_H)$. Tensoring Equation (29) gives $(\Pi_{\mathcal{B}} \Omega_{\mathcal{B}} \Pi_{\mathcal{B}}) \otimes (\Pi_{\mathcal{C}} \Omega_{\mathcal{C}} \Pi_{\mathcal{C}}) \succeq c^2 \Pi$, including for entangled inputs. The two completeness guarantees therefore imply $p_H \geq \mathrm{Tr}[(\Omega_{\mathcal{B}} \otimes \Omega_{\mathcal{C}})\sigma_P] - 2\eta \geq c^2 \,\mathrm{Tr}(\overline{\sigma}_P) - 2\eta - \|\sigma_P - \overline{\sigma}_P\|_1$. Combining this inequality with Equation (28) yields

$$p_H \geq \frac{\gamma^2}{4} p_{\mathrm{str}} - 6\eta - 3\sqrt{2\eta}. \tag{30}$$

Define $\Theta_X := \mathbf{1}[H_X > \gamma/4]$, $\Theta := \Theta_{\mathcal{B}} \otimes \Theta_{\mathcal{C}}$, and $\overline{\sigma}_H := \Theta \sigma_H \Theta$. Soundness of the second pair of threshold measurements gives $\mathrm{Tr}(\overline{\sigma}_H) \geq p_H - 2\eta$, and Section 3.9 gives $\|\sigma_H - \overline{\sigma}_H\|_1 \leq 3\sqrt{2\eta}$.

Recovering the challenge bit. Condition on the event in which Equation (26) holds. Apply Equation (21) with $A = H_X$, $B = (-1)^b \widehat{H}_{X,b}$, $\kappa = \gamma/4$, and $\ell = \gamma/8$. Since $\|A - B\| \leq \gamma/100$, any state supported where $H_X > \gamma/4$ has weight at most

$$\left(\frac{\gamma/100}{\gamma/4 - \gamma/8}\right)^2 = 0.0064 < \frac{1}{100} \tag{31}$$

on the subspace where $(-1)^b \widehat{H}_{X,b} \leq \gamma/8$.

Each party now applies

$$\mathsf{GATI}_{\eta}^{1/2+\gamma/16,\, 1/2-\gamma/16}\big(G_X(\overline{\mathrm{ct}}_X),\, I - G_X(\overline{\mathrm{ct}}_X)\big) \tag{32}$$

and guesses 0 on acceptance and 1 on rejection. If $b = 0$, then outside the subspace bounded in Equation (31), we have $G_X > 1/2 + \gamma/16$, so completeness bounds the rejection probability by the weight of that subspace plus η. If $b = 1$, then outside that subspace, we have $G_X < 1/2 - \gamma/16$, so soundness gives the same bound on the acceptance probability; explicitly, $\mathbf{1}[G_X > 1/2 - \gamma/16] \preceq \mathbf{1}[-\widehat{H}_{X,1} \leq \gamma/8]$. These bounds remain valid when the measured state is entangled with an additional register. Hence, on $\overline{\sigma}_H$, both guesses are correct with probability at least

$$\left(1 - \frac{2}{100}\right) \mathrm{Tr}(\overline{\sigma}_H) - 2\eta. \tag{33}$$

Replacing $\overline{\sigma}_H$ by σ_H costs at most $\|\sigma_H - \overline{\sigma}_H\|_1$, and averaging over the ciphertext lists incurs only the negligible failure probability from Equation (26).

The constructed CORR adversary. The algorithm $\mathcal{A}$ chooses the separate mode of Definition 22, repeats each strong-CPA message pair in all N coordinates, and uses the admissible distribution from Definition 21 that produces a common uniform challenge bit $b_{\mathcal{B}} = b_{\mathcal{C}} = b$. Before it sends the two registers to the receivers, it performs the four threshold measurements above and gives both parties the four outcomes and the same uniform random bit g. If all four measurements accept, each party applies Equation (32) to its ciphertexts; otherwise, both output g.

Let s_H be the probability that all four measurements accept and both final guesses are correct. The preceding bounds and Equation (33) give $s_H \geq (4/5)p_H - 4\eta - 3\sqrt{2\eta} - \mathsf{negl}(\lambda)$, after weakening the constant. When at least one threshold measurement rejects, the shared uniform guess is correct with probability $1/2$. The total success probability is therefore $(1 - p_H)/2 + s_H$, which is at least $1/2 + (3/10)p_H - 4\eta - 3\sqrt{2\eta} - \mathsf{negl}(\lambda)$. After averaging Equation (30), the CORR advantage is at least $(3/40)\varepsilon\gamma^2 - 6\eta - 6\sqrt{2\eta} - \mathsf{negl}(\lambda)$. By Equation (22), this quantity is at least $\varepsilon\gamma^2/20$ for all sufficiently large security parameters in the infinite set. This contradicts SDE-CORR security because $\varepsilon\gamma^2$ is inverse polynomial. □

Remark 2 (A stronger form of strong CPA)**.** *The theorem also holds for a stronger test that accepts whenever both local acceptance eigenvalues are inverse-polynomially far from $1/2$, irrespective of their signs. More precisely, for $X \in \{\mathcal{B}, \mathcal{C}\}$, let $H_X := P_X - I/2$. For an inverse polynomial γ, the joint acceptance projector of this test is $\mathbf{1}[|H_{\mathcal{B}}| \geq \gamma] \otimes \mathbf{1}[|H_{\mathcal{C}}| \geq \gamma]$.*

Decompose this projector according to the four sign choices $(s_{\mathcal{B}}, s_{\mathcal{C}}) \in \{+1, -1\}^2$, where $s_X H_X \geq \gamma$. If the stronger test accepts with non-negligible probability, then one of these four choices occurs with non-negligible probability. The reduction applies the corresponding threshold measurements, sends s_X to party X, and asks that party to use the sign-corrected operator $s_X \widehat{H}_{X,b}$. Since $s_X \widehat{H}_{X,b} \approx (-1)^b s_X H_X$ and $s_X H_X$ is positive on the accepted state, the same argument with N ciphertexts recovers b. When $s_X = -1$, party X simply reverses its final output.

For comparison, let W_+ be the acceptance operator for the test that checks whether the two parties' outputs agree. Then $W_+ = I/2 + 2H_{\mathcal{B}} \otimes H_{\mathcal{C}}$. A positive threshold for $W_+ - I/2$ covers the equal-sign cases $(+1, +1)$ and $(-1, -1)$, whereas a two-sided threshold for $|W_+ - I/2|$ also covers the two opposite-sign cases. Up to an inverse-polynomial change in the threshold, the latter is equivalent to requiring both $|H_{\mathcal{B}}|$ and $|H_{\mathcal{C}}|$ to be inverse-polynomially bounded away from zero. Thus, the stronger test combines ordinary strong CPA, the threshold form of CPA$^+$, and the variants obtained by reversing either party's output; we do not introduce a separate security definition.

Next, we show that the separate mode of CORR-SDE security, used in Theorem 13 is also implied by and hence (by Theorem 13) equivalent to strong CPA security (and equivalently CPA$^+$ security).

Theorem 14 (Separate mode is exactly CPA$^+$)**.** *Let $\mathsf{SDE\text{-}CORR}|_{\mathsf{IND}}$ denote the SDE-CORR game of Definition 22 restricted to adversaries that choose separate mode. Then $\mathsf{SDE\text{-}CORR}|_{\mathsf{IND}}$ security is equivalent to CPA$^+$ security from Definition 18, and hence to strong-CPA security from Definition 19.*

Proof. The proof of Theorem 13 uses only separate mode and therefore proves that $\mathsf{SDE\text{-}CORR}|_{\mathsf{IND}}$ implies strong CPA. Strong CPA and CPA$^+$ are equivalent by Theorem 11. It remains to prove that CPA$^+$ implies SDE-CORR in separate mode.

Fix a separate-mode CORR adversary $\mathcal{A} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ and condition on the public classical information produced by $\mathcal{A}$ before the challenge; the reduction runs $\mathcal{A}$ honestly and finally averages over this information. For $X \in \{\mathcal{B}, \mathcal{C}\}$ and $b \in \{0, 1\}$, let $P_{X,b}$ be the average acceptance operator when X receives independent encryptions of $(\mathrm{m}_{X,1,b}, \ldots, \mathrm{m}_{X,N,b})$ and outputs b. Because the encryptions

for Bob and Charlie use independent randomness, the operator for both parties to answer a fixed pair of challenge bits correctly is $P_{\mathcal{B},b_\mathcal{B}} \otimes P_{\mathcal{C},b_\mathcal{C}}$.

Every joint distribution on $(b_\mathcal{B}, b_\mathcal{C})$ with uniform marginals satisfies $\Pr[00] = \Pr[11] = \theta/2$ and $\Pr[01] = \Pr[10] = (1-\theta)/2$ for some $\theta \in [0,1]$, which may depend on the public key. Define

$$W_{\mathsf{diag}} := \frac{1}{2}\big(P_{\mathcal{B},0} \otimes P_{\mathcal{C},0} + P_{\mathcal{B},1} \otimes P_{\mathcal{C},1}\big),$$
$$W_{\mathsf{off}} := \frac{1}{2}\big(P_{\mathcal{B},0} \otimes P_{\mathcal{C},1} + P_{\mathcal{B},1} \otimes P_{\mathcal{C},0}\big).$$

The CORR acceptance operator is $W_{\mathsf{IND}} = \theta W_{\mathsf{diag}} + (1-\theta) W_{\mathsf{off}}$.

Now use the same state and receivers in CPA$^+$, where the challenge bits are independent. The acceptance operator for the CPA$^+$ condition that the two receivers' correctness bits agree is

$$W_+ := \frac{1}{4} \sum_{b,c \in \{0,1\}} \big(P_{\mathcal{B},b} \otimes P_{\mathcal{C},c} + (I - P_{\mathcal{B},b}) \otimes (I - P_{\mathcal{C},c})\big).$$

Expanding this expression gives

$$W_+ - W_{\mathsf{diag}} = \frac{1}{2}\big((I - P_{\mathcal{B},0}) \otimes (I - P_{\mathcal{C},1}) + (I - P_{\mathcal{B},1}) \otimes (I - P_{\mathcal{C},0})\big) \succeq 0, \tag{34}$$
$$W_+ - W_{\mathsf{off}} = \frac{1}{2}\big((I - P_{\mathcal{B},0}) \otimes (I - P_{\mathcal{C},0}) + (I - P_{\mathcal{B},1}) \otimes (I - P_{\mathcal{C},1})\big) \succeq 0. \tag{35}$$

Thus, $W_{\mathsf{IND}} \preceq W_+$. Every successful separate-mode CORR attack therefore yields a successful CPA$^+$ attack with N challenge pairs. We next reduce the latter attack to the usual game with one pair per receiver.

For a ciphertext list $\overline{\mathrm{ct}}$, let $Z_X(\overline{\mathrm{ct}})$ be the difference between the acceptance operators for outputs 0 and 1. For $j \in \{0, \ldots, N\}$, let $Z_X^{(j)}$ be its expectation when the first j coordinates encrypt their messages indexed by 1 and all later coordinates encrypt their messages indexed by 0. For $j \in [N]$, set $\Delta_{X,j} := Z_X^{(j-1)} - Z_X^{(j)}$. The advantage of the N-pair CPA$^+$ attack over $1/2$ is exactly

$$\frac{1}{8}\operatorname{Tr}\Big[(Z_\mathcal{B}^{(0)} - Z_\mathcal{B}^{(N)}) \otimes (Z_\mathcal{C}^{(0)} - Z_\mathcal{C}^{(N)})\rho\Big] = \frac{1}{8} \sum_{j,k \in [N]} \operatorname{Tr}[(\Delta_{\mathcal{B},j} \otimes \Delta_{\mathcal{C},k})\rho]. \tag{36}$$

Let $T(\lambda)$ be a polynomial upper bound on the number N chosen by the fixed adversary. An ordinary CPA$^+$ adversary runs $\mathcal{A}$ and samples $j, k \xleftarrow{\$} [T]$. If $j > N$ or $k > N$, it submits arbitrary fixed valid message pairs and has both receivers return fixed answers; over the independent challenge bits, the success probability is exactly $1/2$. Otherwise, it submits Bob's coordinate-j pair and Charlie's coordinate-k pair. It gives the receivers copies of pk, N, both message lists, the selected indices, and descriptions of the original receiver algorithms. Each receiver places its one challenge ciphertext in the selected coordinate, independently encrypts the message indexed by 1 in each earlier coordinate and the message indexed by 0 in each later coordinate, and then runs the original receiver on the resulting list. Conditioned on valid j, k, the difference between the receiver's average output operators for challenge bits 0 and 1 is exactly $\Delta_{\mathcal{B},j}$ or $\Delta_{\mathcal{C},k}$. By Equation (36), averaging over j, k gives an ordinary CPA$^+$ advantage equal to $1/T^2$ times the original N-pair advantage. Thus, a non-negligible N-pair advantage contradicts CPA$^+$ security. □

It is not known whether CPA$^+$ security implies full SDE-CORR security. By Theorem 14, separate-mode SDE-CORR is already equivalent to CPA$^+$. Therefore, any separation between full

SDE-CORR and CPA$^+$, or equivalently strong CPA, must use identical mode, in which both parties receive the same ciphertext. Even for $N = 1$, where identical-mode SDE-CORR is IDEN-CPA, this copied-ciphertext setting is not covered by the implications of Kitagawa and Yamakawa [KY25].

In particular, it is not known whether CPA$^+$ implies IDEN-CPA security. The reverse implication is ruled out by the standard two-key construction of [CGLZR23, Theorem 13].

Theorem 15 (IDEN-CPA does not imply CPA$^+$)**.** *Let* Π *be a perfectly correct IDEN-CPA-secure SDE scheme whose message space contains two distinct messages. There is a perfectly correct IDEN-CPA-secure two-key scheme* Π_{2key} *that is not CPA*$^+$ *secure. If* $\Pi.\mathsf{Enc}$ *is classical, then so is* $\Pi_{\mathsf{2key}}.\mathsf{Enc}$*. In particular, this separation holds under post-quantum-secure iO and post-quantum-secure one-way functions.*

Proof. The scheme Π_{2key} generates independent key pairs $(\mathrm{pk}_t, s\mathcal{k}_t) \leftarrow \Pi.\mathsf{GenKey}(1^\lambda)$ for $t \in \{0,1\}$, publishes $(\mathrm{pk}_0, \mathrm{pk}_1)$, and uses $s\mathcal{k}_0 \otimes s\mathcal{k}_1$ as its quantum key. To encrypt m, it samples $t \xleftarrow{\$} \{0,1\}$, computes $\mathrm{ct} \leftarrow \Pi.\mathsf{Enc}(\mathrm{pk}_t, m)$, and outputs (t, ct). Decryption reads t and uses the corresponding key. Perfect correctness is immediate.

Suppose that an adversary $\mathcal{A} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ violates IDEN-CPA security of Π_{2key}. We construct an adversary against Π. Given $(\mathrm{pk}^\star, s\mathcal{k}^\star)$, it samples $J \xleftarrow{\$} \{0,1\}$ and generates an independent key pair. It places the given key pair in position J and the independently generated pair in the other position, and then runs $\mathcal{A}$ for Π_{2key}. It forwards the common message pair output by $\mathcal{A}$ to the challenger for Π and gives each receiver a classical copy of J. Upon receiving the copied ciphertext $\mathrm{ct}^\star$, each receiver runs its original circuit on $(J, \mathrm{ct}^\star)$. The two key pairs are independent and identically distributed, so the ordered public and quantum keys reveal no information about the uniform position J. The simulation is therefore identical to the real IDEN-CPA experiment for Π_{2key}, even if $\mathcal{A}$ jointly processes the two component keys. IDEN-CPA security of Π proves that Π_{2key} is also IDEN-CPA secure.

To violate CPA$^+$ security, fix distinct messages (μ_0, μ_1) and use this pair for both receivers. The algorithm $\mathcal{A}$ gives $s\mathcal{k}_0$ to Bob and $s\mathcal{k}_1$ to Charlie. Bob decrypts when the first ciphertext component is 0 and otherwise returns a private fair bit; Charlie acts analogously when that component is 1. Each receiver is correct with probability $3/4$ and incorrect with probability $1/4$, independently of the other receiver. Their correctness bits therefore agree with probability $(3/4)^2 + (1/4)^2 = 5/8$, so Π_{2key} is not CPA$^+$ secure. Taking Π to be the SDE-CORR-secure scheme of Theorem 22 gives the final assertion in the plain model under post-quantum iO and post-quantum one-way functions. □

The preceding CPA$^+$ attack gives $s\mathcal{k}_0$ to Bob and $s\mathcal{k}_1$ to Charlie, so their joint state is a product state. The next lemma shows that an IDEN-CPA attack whose joint output is a classical mixture of product states cannot separate CPA$^+$ from IDEN-CPA.

Lemma 5 (Separable IDEN-CPA attacks violate strong CPA)**.** *Let* $\mathbf{SDE}$ *be an SDE scheme and let* $\mathcal{A} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ *be an adversary in the IDEN-CPA experiment of Definition 17. Let* u *contain the public key and all classical information output by* $\mathcal{A}$ *before the challenge. For every value* u *that occurs with nonzero probability, let* $\rho^u_{\mathcal{BC}}$ *be the corresponding normalized state on the registers sent to* $\mathcal{B}$ *and* $\mathcal{C}$*. Suppose that this state is separable for every* u*, meaning that* $\rho^u_{\mathcal{BC}} = \sum_z p_{u,z} \rho_{u,z} \otimes \sigma_{u,z}$ *for probabilities* $p_{u,z}$ *and local states* $\rho_{u,z}, \sigma_{u,z}$*. If, for some* $\varepsilon = \varepsilon(\lambda) \in (0, 1/2]$*,*

$$\Pr\left[\mathsf{SDE\text{-}IDEN\text{-}CPA}_{\mathbf{SDE},\mathcal{A}}(1^\lambda) = 1\right] \geq \frac{1}{2} + \varepsilon, \tag{37}$$

then there is an adversary $\mathcal{A}_{\mathsf{str}}$ *in the strong-CPA experiment of Definition 19 such that*

$$\Pr\left[\mathsf{SDE\text{-}STRONG\text{-}CPA}_{\mathbf{SDE},\mathcal{A}_{\mathsf{str}}}(1^\lambda, \varepsilon/4) = 1\right] \geq \frac{\varepsilon^3}{4}. \tag{38}$$

Thus, if **SDE** *is strong-CPA secure, every IDEN-CPA adversary of this form has negligible advantage.*

Proof. The adversary $\mathcal{A}_{\mathsf{str}}$ runs $\mathcal{A}$, uses its common message pair for both strong-CPA decryptors, and equips the two output registers with the circuits of $\mathcal{B}$ and $\mathcal{C}$. We prove that these decryptors satisfy Equation (38).

Claim. Let u be a classical random variable with distribution $(q_u)_u$. For every u, let i be drawn from a distribution $\mathcal{D}_u$, and let $0 \preceq A_{u,i}, B_{u,i} \preceq I$ be operators on two registers. Define $P_u := \mathbb{E}_{i\leftarrow\mathcal{D}_u}[A_{u,i}]$, $Q_u := \mathbb{E}_{i\leftarrow\mathcal{D}_u}[B_{u,i}]$, and $W_u := \mathbb{E}_{i\leftarrow\mathcal{D}_u}[A_{u,i}\otimes B_{u,i}]$. Suppose that $\rho_u = \sum_z p_{u,z}\rho_{u,z}\otimes\sigma_{u,z}$ is separable and $\sum_u q_u \operatorname{Tr}(W_u\rho_u) \geq 1/2+\varepsilon$. For $\Pi_u := \mathbf{1}[P_u \geq 1/2+\varepsilon/4]$ and $\Gamma_u := \mathbf{1}[Q_u \geq 1/2+\varepsilon/4]$, we have

$$\sum_u q_u \operatorname{Tr}\big((\Pi_u\otimes\Gamma_u)\rho_u\big) \geq \frac{\varepsilon^3}{4}. \tag{39}$$

To prove the claim, write $\tau_{u,z} := \rho_{u,z}\otimes\sigma_{u,z}$, and define $a_{u,z} := \operatorname{Tr}(P_u\rho_{u,z})$, $b_{u,z} := \operatorname{Tr}(Q_u\sigma_{u,z})$, and $s_{u,z} := \operatorname{Tr}(W_u\tau_{u,z})$. Since $s_{u,z} = \mathbb{E}_i[\operatorname{Tr}(A_{u,i}\rho_{u,z})\operatorname{Tr}(B_{u,i}\sigma_{u,z})]$ and both factors lie in $[0,1]$, we have $s_{u,z} \leq \min\{a_{u,z}, b_{u,z}\}$. Give each pair (u,z) probability $q_up_{u,z}$, and let G be the set of pairs for which $s_{u,z} \geq 1/2+\varepsilon/2$. Since $0 \leq s_{u,z} \leq 1$, the premise of the claim implies that G has probability at least $\varepsilon/(1-\varepsilon)$.

For every $(u,z)\in G$, both $a_{u,z}$ and $b_{u,z}$ are at least $1/2+\varepsilon/2$. Moreover, $P_u \preceq (1/2+\varepsilon/4)I + (1/2-\varepsilon/4)\Pi_u$, and hence $\operatorname{Tr}(\Pi_u\rho_{u,z}) \geq \varepsilon/(2-\varepsilon)$. The same argument gives $\operatorname{Tr}(\Gamma_u\sigma_{u,z}) \geq \varepsilon/(2-\varepsilon)$. It follows that

$$\sum_u q_u \operatorname{Tr}\big((\Pi_u\otimes\Gamma_u)\rho_u\big) \geq \frac{\varepsilon}{1-\varepsilon}\left(\frac{\varepsilon}{2-\varepsilon}\right)^2 \geq \frac{\varepsilon^3}{4}, \tag{40}$$

where the last inequality uses $0 < \varepsilon \leq 1/2$. This proves the claim.

We apply the claim to the IDEN-CPA attack. Let u contain pk, the common message pair, and all classical information produced before the challenge ciphertext. Conditioned on u, let $i = (b, \mathrm{ct})$, where $b \xleftarrow{\$} \{0,1\}$ and $\mathrm{ct} \leftarrow \mathbf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m}_b)$. Let $A_{u,i}$ be the measurement operator for $\mathcal{B}$ to output b on input ct, and define $B_{u,i}$ analogously for $\mathcal{C}$. Then $\sum_u q_u \operatorname{Tr}(W_u\rho^u_{\mathcal{BC}})$ is exactly the success probability in Equation (37). The operators P_u and Q_u are the two local CPA acceptance operators, so Π_u and Γ_u are exactly the accepting projectors in the strong-CPA experiment with gap $\varepsilon/4$. Thus, Equation (39) proves Equation (38). □

Lemma 5 rules out every separable IDEN-CPA attack arising from a classical mixture over public indices or challenge bits. Thus, separating CPA$^+$ from IDEN-CPA would require an entangled attack that specifically exploits the fact that both parties receive the same ciphertext. Constructing such an attack remains open.

We next compare the two identical-challenge notions.

Theorem 16 (Identical CPA implies identical search [AKL23])**.** *For superlogarithmic message length, identical-CPA security implies identical-search security.*

Proof. Suppose that an identical-search adversary succeeds with probability p. To construct an identical-CPA adversary, first run the search adversary to obtain its two receiver registers and circuits. Independently sample $\mathrm{m}_0, \mathrm{m}_1 \xleftarrow{\$} \mathcal{M}$ and submit $(\mathrm{m}_0, \mathrm{m}_1)$ as the challenge pair. Each CPA receiver runs the corresponding search receiver and outputs 1 exactly when the recovered message equals m_1.

If the challenge bit is 1, both parties are correct with probability p. If it is 0, each CPA receiver first runs its search receiver on the ciphertext and then compares the recovered value

with m_1. The recovered value is independent of the fresh uniform message m_1, so a union bound shows that both parties output 0 with probability at least $1 - 2/|\mathcal{M}|$. The resulting CPA success probability is at least $1/2 + p/2 - 1/|\mathcal{M}|$. If $p \geq 1/|\mathcal{M}| + \varepsilon$ for non-negligible ε, this lower bound is $1/2 + \varepsilon/2 - 1/(2|\mathcal{M}|)$. The final term is negligible for superlogarithmic message length, contradicting identical-CPA security. □

Identical-challenge search and CPA security do not imply their independent-challenge counterparts; the corresponding two-key constructions appear in [CGLZR23, Theorems 12 and 13]. The CPA construction also fails independent-search security: assigning one component key to each receiver allows both to decrypt whenever their independently sampled ciphertext indices identify the assigned keys, an event of constant probability. We omit this additional separation between identical and independent challenges from the figure for readability.

Search security does not imply CPA security either. Given a search-secure scheme over $\mathcal{M}_0$, enlarge its message space to $\{0,1\} \times \mathcal{M}_0$ and include the first component in the clear in every ciphertext. Exact recovery still requires recovering the $\mathcal{M}_0$ component, and superlogarithmic message length makes the resulting change in the uniform-guessing probability negligible. By contrast, messages that differ only in the first component are perfectly distinguishable. This construction gives the search-to-CPA separations in the figure, including the analogous separations between identical- and independent-challenge notions.

It remains to prove the separations with IND-CPA on the left. One scheme suffices: if it is IND-CPA secure but not IDEN-SEARCH secure, then the implications above show that it is neither IDEN-CPA nor IND-SEARCH secure. It is also neither CPA$^+$ nor strong-CPA secure, since both imply IND-SEARCH. The next construction, due to Coladangelo, Liu, Liu, and Zhandry, uses quantum states associated with shifted subspaces and a logarithmic-length public index in each ciphertext. The proof must also accommodate the unrelated message pairs allowed for the two receivers in IND-CPA.

Theorem 17 (Independent CPA does not imply strong CPA)**.** *Assuming post-quantum-secure indistinguishability obfuscation and post-quantum-secure one-way functions, there is a perfectly correct SDE scheme that is* SDE-IND-CPA *secure but not* SDE-STRONG-CPA *secure. The same scheme is not* SDE-IDEN-SEARCH *secure and establishes every displayed separation in Figure 1 whose left-hand notion is* SDE-IND-CPA*.*

Proof. Let $n = 2\lambda$, let $L = \lceil c \log \lambda \rceil$ for any fixed constant $c > 0$, set $R := 2^L$, and take $\mathcal{M} = \{0,1\}^\lambda$. For a subspace $A \leq \mathbb{F}_2^n$ and shifts $s, t \in \mathbb{F}_2^n$, define the coset state

$$|A_{s,t}\rangle := \frac{1}{\sqrt{|A|}} \sum_{a \in A} (-1)^{\langle a,t \rangle} |a + s\rangle .$$

For all sufficiently large λ, we have $R \geq 4$. For each $i \in [L]$, key generation independently samples an $n/2$-dimensional subspace $A_i \leq \mathbb{F}_2^n$ and shifts $s_i, t_i \xleftarrow{\$} \mathbb{F}_2^n$. It publishes obfuscated programs that test membership in $A_i + s_i$ and $A_i^\perp + t_i$, and it sets $\mathsf{sk} := \bigotimes_{i=1}^L |(A_i)_{s_i,t_i}\rangle$.

For $a \in \{0,1\}$, define $B_{i,0} := A_i$, $B_{i,1} := A_i^\perp$, $c_{i,0} := s_i$, and $c_{i,1} := t_i$. For a public index $r \in \{0,1\}^L$, let $B_r := \prod_i B_{i,r_i}$, let $c_r := (c_{1,r_1}, \ldots, c_{L,r_L})$, and set $S_r := B_r + c_r$. The public programs determine membership in S_r. For $m \in \mathcal{M}$, define

$$P_{r,m}(z) := \begin{cases} m, & z \in S_r, \\ \perp, & z \notin S_r. \end{cases}$$

Encryption samples $r \xleftarrow{\$} \{0,1\}^L$ and outputs $\mathrm{ct} = (r, \mathsf{iO}(P_{r,m}))$. To decrypt, apply a Hadamard transform to the register containing $|(A_i)_{s_i,t_i}\rangle$ exactly when $r_i = 1$, evaluate the obfuscated program reversibly, and undo these operations. The transformed key state is supported on S_r, so correctness is perfect. This is the SDE construction of [CLLZ21], with the program enlarged to output λ bits.

We use the following known no-cloning property of coset states. Given the public membership programs and the tensor product of coset states above, no QPT algorithm $\mathcal{A}$ that distributes two registers can enable two noncommunicating parties, after descriptions of the A_i are revealed, to output (r, z) and (r', z') satisfying

$$r \neq r',\ z \in S_r,\ \text{and } z' \in S_{r'}. \tag{41}$$

The computational statement follows from [CLLZ21, Theorem 4.18]; the information-theoretic claim assumed there was later proved by Culf and Vidick [CV22]. The one-coordinate game in the cited theorem asks one party for a point in $A_J + s_J$ and the other for a point in $A_J^\perp + t_J$. To reduce the statement above for two distinct indices to that game, place its challenge in a uniformly hidden coordinate J and generate every other coordinate honestly. The reduction neither reveals nor uses J while it runs $\mathcal{A}$, performs the threshold measurements, or chooses the two indices. Because the coordinates are independently and identically distributed, conditioned on $r \neq r'$, the probability that $r_J \neq r'_J$ is $\mathrm{dist}(r, r')/L \geq 1/L$. After r and r' have been chosen, but before the two registers are sent, the reduction sends the register whose Jth index bit is 0 to the party required to output a point in $A_J + s_J$, and it sends the other register to the party required to output a point in $A_J^\perp + t_J$. The reduction loses a factor of at most L.

The scheme is not IDEN-SEARCH secure. Fix $r^\star \in \{0,1\}^L$, apply Hadamard transforms according to $r^\star$, and measure the quantum key to obtain $z^\star \in S_{r^\star}$. Give both parties a classical copy of $z^\star$ and the same uniform value $g \in \mathcal{M}$. On a common ciphertext whose first component is r, both parties recover the message if $r = r^\star$, and otherwise both output g. Their joint success probability is

$$\frac{1}{R} + \left(1 - \frac{1}{R}\right)\frac{1}{|\mathcal{M}|} = \frac{1}{|\mathcal{M}|} + \frac{1}{R}\left(1 - \frac{1}{|\mathcal{M}|}\right), \tag{42}$$

which exceeds the uniform-guessing probability by an inverse polynomial.

We now prove IND-CPA security in the generalized game in which the parties may choose unrelated message pairs. Suppose, toward a contradiction, that an adversary $\mathcal{A} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ succeeds with probability at least $1/2+\varepsilon$ for inverse-polynomial ε. The adversary's classical message-pair and circuit-output registers remain part of the state, and all operations below are controlled by them; we suppress these registers. For $X \in \{\mathcal{B}, \mathcal{C}\}$ and $r \in \{0,1\}^L$, let $P_{X,r}$ be X's acceptance operator when the first ciphertext component is fixed to r, the challenge bit is uniform, and encryption is otherwise honest. Set $P_X := R^{-1}\sum_{r\in\{0,1\}^L} P_{X,r}$. The two parties receive independent challenge bits, public indices, and encryptions, so their joint acceptance operator is $P_{\mathcal{B}} \otimes P_{\mathcal{C}}$. Hence, the state ρ distributed by $\mathcal{A}$ satisfies

$$\mathrm{Tr}[(P_{\mathcal{B}} \otimes P_{\mathcal{C}})\rho] \geq \frac{1}{2} + \varepsilon. \tag{43}$$

If either party chooses the same message twice, its challenge bit is independent of its view and its acceptance operator is $I/2$, which makes the joint success probability at most $1/2$. Thus, both message pairs are distinct.

We first find distinct public indices for which the two receivers have nontrivial prediction probabilities. Every threshold measurement below is the measurement from Lemma 3, whose guarantees remain valid when the measured register is entangled with another register. Choose the inverse-polynomial error parameters sufficiently small, as fixed polynomials in ε/L, that all resulting errors

are absorbed by the inverse-polynomial lower bounds below. These measurements have reversible implementations using the public encryption and receiver circuits while preserving their randomness and work registers; no cryptographic assumption is used in this step. The algorithm $\mathcal{A}$ performs every threshold measurement and chooses both indices before distributing the two registers. It uses only the public key, the message pairs, and the public encryption algorithm. Descriptions of the hidden subspaces are revealed only later, when the two receivers run the point-producing algorithms described below.

Apply the threshold measurement to $P_{\mathcal{B}} \otimes P_{\mathcal{C}}$, with completeness threshold $1/2 + 3\varepsilon/4$ and soundness threshold $1/2 + 2\varepsilon/3$. The inequality $M \preceq \tau I + (1-\tau)\mathbf{1}[M \geq \tau]$ and Equation (43) imply that this measurement accepts with probability $\Omega(\varepsilon)$. Apart from the chosen approximation error, every pair of local eigenvalues $(\lambda_{\mathcal{B}}, \lambda_{\mathcal{C}})$ in the accepted state satisfies

$$\lambda_{\mathcal{B}}\lambda_{\mathcal{C}} > \frac{1}{2} + \frac{2\varepsilon}{3}. \tag{44}$$

Both eigenvalues exceed $1/2 + 2\varepsilon/3$, and at least one exceeds $1/\sqrt{2} > 11/16$. Choose one party uniformly, denote it by H, and denote the other by W; the letter H merely identifies the party tested against the higher constant threshold. Apply the threshold measurement to P_H with completeness threshold $11/16$ and soundness threshold $43/64$. With constant probability, H is a party whose eigenvalue exceeds $11/16$, and the measurement accepts. Apart from the chosen approximation error, the resulting state is supported on local spectral subspaces on which

$$P_H \succeq \frac{43}{64} I \text{ and } P_W \succeq \left(\frac{1}{2} + \frac{2\varepsilon}{3}\right) I. \tag{45}$$

Here and below, such inequalities are understood after compression to the stated post-measurement support, and I denotes the identity on that support.

Sample $r_W \xleftarrow{\$} \{0,1\}^L$ and apply the threshold measurement to P_{W,r_W} with completeness threshold $1/2 + \varepsilon/4$ and soundness threshold $1/2 + 7\varepsilon/32$. Averaging the threshold bound over r_W and using Equation (45) gives acceptance probability $\Omega(\varepsilon)$. Next sample $r_H \xleftarrow{\$} \{0,1\}^L \setminus \{r_W\}$. On the accepted support, $P_H \succeq (43/64)I$, and $P_{H,r_W} \preceq I$; hence

$$\frac{1}{R-1}\sum_{r \neq r_W} P_{H,r} = \frac{RP_H - P_{H,r_W}}{R-1} \succeq \frac{R(43/64) - 1}{R-1} I \succeq \frac{9}{16} I, \tag{46}$$

where the last inequality uses $R \geq 4$. Applying the threshold measurement to P_{H,r_H} with completeness threshold $17/32$ and soundness threshold $67/128$ therefore accepts with constant probability. A measurement on one party does not disturb spectral support already established on the other party's register. Thus, before the two registers are distributed, with probability $\Omega(\varepsilon^2)$ we obtain distinct indices $r_W \neq r_H$ and a state that, apart from the chosen approximation error, is supported where

$$P_{W,r_W} > \frac{1}{2} + \frac{7\varepsilon}{32} \text{ and } P_{H,r_H} > \frac{1}{2} + \frac{3}{128}. \tag{47}$$

It remains to show that a decryptor that predicts the challenge bit for a fixed public index r can produce a point in S_r, even for an arbitrary local message pair $(\mu_{X,0}, \mu_{X,1})$. Let $\Delta_X := \mu_{X,0} \oplus \mu_{X,1}$, where $a \cdot \Delta_X$ denotes Δ_X when $a = 1$ and the all-zero string when $a = 0$. For $q \xleftarrow{\$} B_r^{\perp}$ and $d \xleftarrow{\$} \{0,1\}$, define

$$Q_{X,r,q,d}(z) := \begin{cases} \mu_{X,0} \oplus (d \oplus \langle q, z\rangle) \cdot \Delta_X, & z \in S_r, \\ \perp, & z \notin S_r. \end{cases} \tag{48}$$

For $z \in S_r = B_r + c_r$, we have $\langle q, z\rangle = \langle q, c_r\rangle$. Thus, for $b = d \oplus \langle q, c_r\rangle$, the circuits satisfy $Q_{X,r,q,d} \equiv P_{r,\mu_{X,b}}$. Moreover, b is uniform because d is uniform. By iO security from Definition 4, replacing the honest fixed-r program by $Q_{X,r,q,d}$ changes the prediction probability by at most a negligible amount. The following claim states the required conclusion when the receiver may be entangled with an additional register.

Fixed-index recovery claim. Let τ be an efficiently preparable subnormalized state, that is, a positive operator satisfying $\mathrm{Tr}(\tau) \le 1$, that contains the honestly generated classical setup, a public index r, a distinct message pair (μ_0, μ_1), a receiver register, and an arbitrary additional register. Suppose that $p := \mathrm{Tr}(\tau)$ is inverse polynomial and that, in the honest fixed-r challenge-bit experiment, the receiver is correct with total probability at least $(1/2 + \delta)p$, where δ is inverse polynomial. After descriptions of the A_i are revealed, there is a local QPT algorithm that leaves the additional register unchanged and outputs $z \in S_r$ with probability $\Omega(p\delta^2)$.

To prove the claim, an adversary in the iO security experiment of Definition 4 generates the honest setup and prepares τ, and thus knows c_r for this comparison. It samples $q \xleftarrow{\$} B_r^{\perp}$ and $d \xleftarrow{\$} \{0,1\}$, sets $b = d \oplus \langle q, c_r\rangle$, and gives its iO challenger equally padded descriptions of P_{r,μ_b} and $Q_{r,q,d}$, which compute the same function. It runs the receiver on the returned obfuscation and tests whether the output is b; if preparation of τ fails, it instead outputs a fair bit. With P_{r,μ_b}, this is exactly the honest fixed-r test for uniform b. With $Q_{r,q,d}$, XORing the receiver's output with d predicts $\langle q, c_r\rangle$. Post-quantum iO therefore preserves the total prediction advantage on τ up to $\mathsf{negl}(\lambda)$, or up to $\mathsf{negl}(\lambda)/p$ after conditioning on successful preparation of τ. The argument works directly with τ; it neither prepares the normalized state nor requires inverse-polynomial probability for each individual value of r.

After the A_i are revealed, let G_r be a full-rank matrix whose rows span $B_r^{\perp}$, and write $q = G_r^{\mathsf{T}} y$. The predictor now guesses $\langle q, c_r\rangle = \langle y, G_r c_r\rangle$. The quantum Goldreich–Levin lemma [CLLZ21, Lemma B.12] allows the predictor to hold an arbitrary quantum register and recovers $G_r c_r$ with probability $\Omega(\delta^2)$, conditioned on successful preparation of τ. More precisely, let Γ_Q be twice the prediction advantage when using Q, weighted by the trace p of τ. Then iO security and the averaging identity in the cited lemma give $\Gamma_Q \ge 2\delta p - \mathsf{negl}(\lambda)$, and Cauchy–Schwarz gives $\Pr[\mathsf{recover}] \ge \Gamma_Q^2/p = \Omega(p\delta^2)$. These bounds average over the honest setup and the public index. Solving $G_r z = G_r c_r$ yields $z \in B_r + c_r = S_r$.

The algorithm producing z computes q, the program in Equation (48), and the classical obfuscator reversibly, keeping every random coin in a work register. The iO comparison samples q classically and never invokes iO on a superposition of circuits. The prediction-to-recovery guarantee of the cited Goldreich–Levin lemma remains valid when the receiver register is entangled with the additional register. This proves the claim.

Apply the claim first to W, treating H's register as the additional register that must remain unchanged. The first inequality in Equation (47) gives conditional probability $\Omega(\varepsilon^2)$ of producing a valid point. Verify the output publicly and condition on successful verification, thereby obtaining a subnormalized state. Because this operation acts only on W, the second spectral-support condition in Equation (47) remains valid up to the chosen measurement error; see [CLLZ21, Claim 6.18]. The subnormalized state still has inverse-polynomial trace, so applying the claim to H gives constant conditional probability of producing a valid point for r_H. The two local algorithms may run concurrently; we describe them sequentially only to analyze their joint success. Including the $\Omega(\varepsilon^2)$ probability of reaching the state in Equation (47), the algorithms produce valid points for the distinct indices r_W and r_H with probability $\Omega(\varepsilon^4)$ after accounting for those errors. The additional $1/L$ loss in the one-coordinate reduction is inverse polynomial, contradicting Equation (41). Thus, the scheme is IND-CPA secure.

Together with Equation (42), this result separates IND-CPA from IDEN-SEARCH. Since IDEN-CPA implies IDEN-SEARCH and IND-SEARCH implies IDEN-SEARCH, it also gives the corresponding two separations. Finally, if the scheme were strong-CPA secure, then Theorem 11 would imply CPA$^+$ security, then IND-SEARCH security, and finally IDEN-SEARCH security, contradicting Equation (42). Hence, the scheme is not strong-CPA secure and, by the same equivalence, is not CPA$^+$ secure. □

4.3.1 Summary of the Main Relations

For convenience, we summarize the main relations established above. As in Section 4.3, every statement involving search security assumes superlogarithmic message length.

1. SDE-CORR security implies strong-CPA security (and hence CPA$^+$ security by [KY25]) by Theorem 13, and it directly implies IND-CPA and IDEN-CPA security by Lemma 4.

2. The restriction of SDE-CORR to separate mode is equivalent to CPA$^+$ security and hence to strong-CPA security, even with polynomially many challenge-message pairs and every admissible challenge sampler; see Theorem 14.

3. Strong-CPA security, or equivalently CPA$^+$ security, implies IND-CPA and IND-SEARCH security, while IND-SEARCH and strong-search security are equivalent and imply IDEN-SEARCH security; see Theorem 11.

4. Strong-CPA security strictly implies strong-search security under the assumptions stated in Theorem 12. Moreover, IDEN-CPA security implies IDEN-SEARCH security by Theorem 16.

5. IDEN-CPA security does not imply CPA$^+$ security by Theorem 15. Assuming post-quantum-secure iO and post-quantum-secure one-way functions, IND-CPA security does not imply IDEN-CPA, IND-SEARCH, IDEN-SEARCH, CPA$^+$, or strong-CPA security; all these separations are witnessed by the scheme in Theorem 17.

6. IDEN-SEARCH security does not imply IND-SEARCH security, and IDEN-CPA security implies neither IND-CPA nor IND-SEARCH security [CGLZR23, Theorems 12 and 13]. Moreover, placing one component of every message in the clear shows that IND-SEARCH implies neither IND-CPA nor IDEN-CPA, and that IDEN-SEARCH implies neither IDEN-CPA nor IND-CPA.

5 Unclonable Puncturable Obfuscation Security Definitions and Copy-Protection of General Functionalities

In this section, we present security definitions for unclonable puncturable obfuscation and copy-protection of general functionalities.

5.1 Unclonable Puncturable Obfuscation definitions

First, we recall the syntax and correctness of **UPO** from [AB24, ABH+26], then in Section 5.1.1 we recall the prior security definitions. Then, in Section 5.1.2 we present our new definition and show that our new definition implies all prior UPO security notions considered in previous works, except UPO plus. Let $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda\in\mathbb{N}}$ be a keyed circuit class, where $\mathcal{C}_\lambda = \{C_k : \mathcal{X}_\lambda \to \mathcal{Y}_\lambda\}_{k\in\mathcal{K}_\lambda}$, $\mathcal{X}_\lambda := \{0,1\}^{\ell_{\mathsf{in}}(\lambda)}$, and $\mathcal{Y}_\lambda := \{0,1\}^{\ell_{\mathsf{out}}(\lambda)}$.

Definition 23 (Unclonable puncturable obfuscation)**.** *A UPO scheme* $\Pi = (\mathsf{Obf}, \mathsf{Eval})$ *for* $\mathcal{C}$ *consists of a pair of efficient algorithms as follows.*

1. $\rho_k \leftarrow \mathsf{Obf}(1^\lambda, C_k)$ *: is a QPT algorithm that takes as input a circuit* $C_k \in \mathcal{C}_\lambda$ *and outputs a quantum obfuscation* ρ_k*.*
2. $(\rho'_k, y) \leftarrow \mathsf{Eval}(\rho_k, x)$ *: is a QPT algorithm that takes as input a quantum obfuscation* ρ_k *and an input* $x \in \mathcal{X}_\lambda$*, and outputs an updated quantum obfuscation* ρ'_k *and* $y \in \mathcal{Y}_\lambda$*.*

Correctness *We require that there exists a negligible function* $\mathsf{negl}(\cdot)$ *such that, for every* $\lambda \in \mathbb{N}$*, every* $k \in \mathcal{K}_\lambda$*, and every* $x \in \mathcal{X}_\lambda$*,*

$$\Pr\Big[y = C_k(x) \mid \rho_k \leftarrow \mathsf{Obf}(1^\lambda, C_k), (\rho'_k, y) \leftarrow \mathsf{Eval}(\rho_k, x)\Big] \geq 1 - \mathsf{negl}(\lambda). \tag{49}$$

As in [AB24, Remark 8], the almost-as-good-as-new lemma and a quantum union bound give polynomial reusability. We say that the scheme is perfectly correct when the probability above is one.

Throughout this section, whenever a UPO security experiment obfuscates a circuit produced by one of its specified puncturing, programming, or compilation algorithms, the UPO scheme and its correctness requirement are understood to apply to that circuit as well.

5.1.1 Prior Definitions:Generalized UPO Security

We next recall generalized puncturing from [AB24, ABH+26]. A keyed circuit class $\mathcal{C}$ is generalized puncturable if there exists a deterministic PPT algorithm $\mathsf{GenPuncture}$ with the following syntax.

Definition 24 (Generalized puncturing)**.** $G^\star_{k,x_\mathcal{B},x_\mathcal{C},\mu_\mathcal{B},\mu_\mathcal{C}} \leftarrow \mathsf{GenPuncture}(k, x_\mathcal{B}, x_\mathcal{C}, \mu_\mathcal{B}, \mu_\mathcal{C})$ *: takes as input a key* $k \in \mathcal{K}_\lambda$*, points* $x_\mathcal{B}, x_\mathcal{C} \in \mathcal{X}_\lambda$*, and circuits* $\mu_\mathcal{B}, \mu_\mathcal{C} : \mathcal{X}_\lambda \to \mathcal{Y}_\lambda$*, and outputs a circuit* $G^\star_{k,x_\mathcal{B},x_\mathcal{C},\mu_\mathcal{B},\mu_\mathcal{C}}$ *satisfying*

$$G^\star_{k,x_\mathcal{B},x_\mathcal{C},\mu_\mathcal{B},\mu_\mathcal{C}}(x) := \begin{cases} \mu_\mathcal{B}(x), & x = x_\mathcal{B}, \\ \mu_\mathcal{C}(x), & x = x_\mathcal{C} \text{ and } x_\mathcal{C} \neq x_\mathcal{B}, \\ C_k(x), & \text{otherwise.} \end{cases} \tag{50}$$

All compared circuits are padded to the same size.

For any distribution $\mathcal{D}_X$ over $\mathcal{X}_\lambda \times \mathcal{X}_\lambda$ and any QPT adversarial triplet $\mathcal{A}_{\mathsf{UPO}} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$, consider the following experiment.
$\mathsf{GenUPO}_{\Pi,\mathcal{A}_{\mathsf{UPO}},\mathcal{D}_{\mathcal{X}},\mathcal{C}}(1^\lambda)$:

1. $\mathcal{A}(1^\lambda)$ sends $(k, \mu_{\mathcal{B}}, \mu_{\mathcal{C}})$ to Ch, where $k \in \mathcal{K}_\lambda$ and $\mu_{\mathcal{B}}, \mu_{\mathcal{C}} : \mathcal{X}_\lambda \to \mathcal{Y}_\lambda$.
2. Ch samples $(x_{\mathcal{B}}, x_{\mathcal{C}}) \leftarrow \mathcal{D}_X(1^\lambda)$ and $b \xleftarrow{\$} \{0,1\}$, and computes $G^\star \leftarrow \mathsf{GenPuncture}(k, x_{\mathcal{B}}, x_{\mathcal{C}}, \mu_{\mathcal{B}}, \mu_{\mathcal{C}})$.
3. Ch generates $\rho_b \leftarrow \mathsf{Obf}(1^\lambda, C_k)$ if $b = 0$, and $\rho_b \leftarrow \mathsf{Obf}(1^\lambda, G^\star)$ if $b = 1$, and sends ρ_b to $\mathcal{A}$.
4. $\mathcal{A}$ outputs a joint state $\sigma_{\mathsf{R}_1,\mathsf{R}_2}$ and sends register R_1 to $\mathcal{B}$ and register R_2 to $\mathcal{C}$.
5. After the split, Ch sends $x_{\mathcal{B}}$ to $\mathcal{B}$ and $x_{\mathcal{C}}$ to $\mathcal{C}$.
6. $\mathcal{B}$ and $\mathcal{C}$ output $\widehat{b_{\mathcal{B}}}, \widehat{b_{\mathcal{C}}}$.
7. The output of the experiment is 1 if $\widehat{b_{\mathcal{B}}} = \widehat{b_{\mathcal{C}}} = b$.

Definition 25 ($\mathcal{D}_X$-generalized UPO security)**.** *A UPO scheme Π satisfies $\mathcal{D}_X$-generalized UPO security with respect to $\mathcal{C}$ if, for every QPT adversarial triplet $\mathcal{A}_{\mathsf{UPO}} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$, there exists a negligible function $\mathsf{negl}(\cdot)$ such that*

$$\Pr\left[\mathsf{GenUPO}_{\Pi,\mathcal{A}_{\mathsf{UPO}},\mathcal{D}_{\mathcal{X}},\mathcal{C}}(1^\lambda) = 1\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda). \tag{51}$$

We use the following standard instantiations of the challenge distribution.

Remark 3. *For a distribution $\mathcal{D}$ over $\mathcal{X}_\lambda$, let $\mathcal{D}^2$ denote the product distribution obtained by sampling $x_{\mathcal{B}}, x_{\mathcal{C}} \leftarrow \mathcal{D}$ independently. Let $\mathsf{ID}_{\mathcal{D}}$ denote the diagonal distribution obtained by sampling $x \leftarrow \mathcal{D}$ and outputting (x, x). Let U denote the uniform distribution on $\mathcal{X}_\lambda^2$. In particular, $\mathsf{ID}_{\mathsf{U}} := \mathsf{ID}_{\mathsf{U}}$ denotes the identical-uniform distribution, which samples $x \xleftarrow{\$} \mathcal{X}_\lambda$ and outputs (x, x). Thus U-generalized UPO is the independent-uniform or product-challenge notion, while $\mathsf{ID}_{\mathcal{D}}$-generalized UPO is the diagonal notion and ID_{U}-generalized UPO is the identical-uniform notion.*

5.1.2 Correlated-Challenge Security

We generalize the prior definition of UPO in three ways: first, we allow the challenge bits for $\mathcal{B}$ and $\mathcal{C}$ to be arbitrarily correlated; next, we allow both the challenge points and the two puncturing bits to be sampled jointly, thus allowing them to be correlated; finally, we allow the sampler to output classical auxiliary information z_{aux} that is given to $\mathcal{A}$ before she chooses the circuit and classical auxiliary information $w_{\mathcal{B}}, w_{\mathcal{C}}$ that is given to $\mathcal{B}, \mathcal{C}$, respectively, after the split[7]. We require the challenge bits to have uniform marginals conditioned on all the auxiliary information and the challenge points, and require each challenge point to have high average conditional min-entropy, as defined in Definition 6, conditioned on the pre-split auxiliary information.[8]

[7] We note that auxiliary information leakage is a natural generalization which can become crucial for some of the applications of UPO (in fact, this generalization is crucial for the application to copy protection for compute-and-compare functions, see Theorem 26).

[8] Arbitrary correlation without these restrictions is not tolerable, since the auxiliary information or the challenge points may trivially reveal a challenge bit.

Definition 26 (Correlated High min-entropy UPO sampler)**.** *Let $\kappa := \kappa(\lambda) = \lambda^c$ for some constant $c > 0$. A correlated UPO sampler with κ-entropic marginals for a keyed-circuit class $\mathcal{C} = \{C_\lambda\}_{\lambda\in\mathbb{N}}$ with input space $\mathcal{X}$ is a QPT algorithm* Samp *that, on input 1^λ, outputs $(z_{\mathsf{aux}}, x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}, w_{\mathcal{B}}, w_{\mathcal{C}}) \leftarrow \mathsf{Samp}(1^\lambda)$, where $z_{\mathsf{aux}} \in \mathcal{Z}_\lambda^{\mathsf{aux}}$ is a polynomial-length classical auxiliary string, $x_{\mathcal{B}}, x_{\mathcal{C}} \in \mathcal{X}_\lambda$ and $b_{\mathcal{B}}, b_{\mathcal{C}} \in \{0,1\}$, and $w_{\mathcal{B}}, w_{\mathcal{C}}$ are polynomial-length classical auxiliary strings, such that*

1. *For every $(\bar{z}_{\mathsf{aux}}, \bar{x}_{\mathcal{B}}, \bar{x}_{\mathcal{C}}, \bar{w}_{\mathcal{B}}, \bar{w}_{\mathcal{C}})$ in the support of $(z_{\mathsf{aux}}, x_{\mathcal{B}}, x_{\mathcal{C}}, w_{\mathcal{B}}, w_{\mathcal{C}})$, every $X \in \{\mathcal{B}, \mathcal{C}\}$, and every $a \in \{0,1\}$,*
$$\Pr\left[b_X = a \mid \begin{smallmatrix} z_{\mathsf{aux}}=\bar{z}_{\mathsf{aux}},\ x_{\mathcal{B}}=\bar{x}_{\mathcal{B}},\ x_{\mathcal{C}}=\bar{x}_{\mathcal{C}}, \\ w_{\mathcal{B}}=\bar{w}_{\mathcal{B}},\ w_{\mathcal{C}}=\bar{w}_{\mathcal{C}} \end{smallmatrix}\right] = \frac{1}{2}. \tag{52}$$
2. *For every $X \in \{\mathcal{B}, \mathcal{C}\}$, $\widetilde{H}_\infty(x_X \mid z_{\mathsf{aux}}) \geq \kappa(\lambda)$.*

We identify a sampler with no explicit auxiliary output with the special case $z_{\mathsf{aux}} := w_{\mathcal{B}} := w_{\mathcal{C}} := \bot$.

For $k, x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}, \mu_{\mathcal{B}}, \mu_{\mathcal{C}}$ as above, define

$$G^{\mathsf{ac}}_{k,x_{\mathcal{B}},x_{\mathcal{C}},b_{\mathcal{B}},b_{\mathcal{C}},\mu_{\mathcal{B}},\mu_{\mathcal{C}}}(x) := \begin{cases} \mu_{\mathcal{B}}(x), & x = x_{\mathcal{B}},\ b_{\mathcal{B}} = 1, \\ \mu_{\mathcal{C}}(x), & x = x_{\mathcal{C}},\ x_{\mathcal{C}} \neq x_{\mathcal{B}},\ b_{\mathcal{C}} = 1, \\ C_k(x), & \text{otherwise.} \end{cases} \tag{53}$$

Equivalently, $G^{\mathsf{ac}} \leftarrow \mathsf{Correlated\text{-}Puncture}(k, x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}, \mu_{\mathcal{B}}, \mu_{\mathcal{C}})$. If $x_{\mathcal{B}} = x_{\mathcal{C}}$ and $b_{\mathcal{B}} \neq b_{\mathcal{C}}$, we follow the convention used in [ABH+26] that $\mathcal{B}$'s challenge bit takes priority, i.e., the common point $x_{\mathcal{B}} = x_{\mathcal{C}}$ is punctured using $\mu_{\mathcal{B}}$ if $b_{\mathcal{B}} = 1$, and is left unpunctured if $b_{\mathcal{B}} = 0$, oblivious of $\mathcal{C}$'s challenge bit $b_{\mathcal{C}}$.

Definition 27 (Correlated UPO security)**.** *For any QPT adversarial triplet $\mathcal{A}_{\mathsf{UPO}} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ against a UPO scheme $\Pi = (\mathsf{Obf}, \mathsf{Eval})$ and any correlated UPO sampler* Samp *with κ-entropic marginals, write $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$, where $\mathcal{A}_0$ may retain an arbitrary quantum state for $\mathcal{A}_1$, and consider the following experiment.*
Correlated-UPO$_{\Pi,\mathcal{A}_{\mathsf{UPO}},\mathsf{Samp},\mathcal{C}}(1^\lambda)$:

1. *Ch samples*
$$(z_{\mathsf{aux}}, x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}, w_{\mathcal{B}}, w_{\mathcal{C}}) \leftarrow \mathsf{Samp}(1^\lambda)$$
and sends z_{aux} to $\mathcal{A}_0$.
2. *$\mathcal{A}_0(1^\lambda, z_{\mathsf{aux}})$ sends $(k, \mu_{\mathcal{B}}, \mu_{\mathcal{C}})$ to Ch, where $k \in \mathcal{K}_\lambda$ and $\mu_{\mathcal{B}}, \mu_{\mathcal{C}} : \mathcal{X}_\lambda \to \mathcal{Y}_\lambda$, while retaining an arbitrary quantum state.*
3. *Ch computes $G^{\mathsf{ac}} \leftarrow$ Correlated-Puncture$(k, x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}, \mu_{\mathcal{B}}, \mu_{\mathcal{C}})$, generates $\rho \leftarrow \mathsf{Obf}(1^\lambda, G^{\mathsf{ac}})$, and sends ρ to $\mathcal{A}_1$.*
4. *$\mathcal{A}_1$ outputs a joint state $\sigma_{\mathsf{R}_1,\mathsf{R}_2}$ and sends register R_1 to $\mathcal{B}$ and register R_2 to $\mathcal{C}$.*
5. *After the split, Ch sends $(x_{\mathcal{B}}, w_{\mathcal{B}})$ to $\mathcal{B}$ and $(x_{\mathcal{C}}, w_{\mathcal{C}})$ to $\mathcal{C}$.*
6. *$\mathcal{B}$ and $\mathcal{C}$ output $\widehat{b}_{\mathcal{B}}, \widehat{b}_{\mathcal{C}}$.*
7. *The output of the experiment is 1 if $\widehat{b}_{\mathcal{B}} = b_{\mathcal{B}}$ and $\widehat{b}_{\mathcal{C}} = b_{\mathcal{C}}$.*

A UPO scheme Π *satisfies* κ*-entropic* Correlated-UPO *security with respect to* $\mathcal{C}$ *if, for every correlated UPO sampler* Samp *with* κ*-entropic marginals for* $\mathcal{C}$*, and every QPT adversarial triplet* $\mathcal{A}_{\mathsf{UPO}}$*, there exists a negligible function* $\mathsf{negl}(\cdot)$ *such that*

$$\Pr\Big[\textit{Correlated-UPO}_{\Pi,\mathcal{A}_{\mathsf{UPO}},\mathsf{Samp},\mathcal{C}}(1^\lambda)=1\Big]\le\frac{1}{2}+\mathsf{negl}(\lambda). \tag{54}$$

Definition 28 ($i\mathcal{O}$-security for UPO, adapted from [ABH+26, Definition 3.15])**.** *A UPO scheme* $\Pi=(\mathsf{Obf},\mathsf{Eval})$ *for a circuit class* $\mathcal{C}$ *satisfies* $i\mathcal{O}$*-security if, for every QPT adversarial pair* $\mathcal{A}_{i\mathcal{O}}=(\mathcal{A}_{i\mathcal{O},0},\mathcal{A}_{i\mathcal{O},1})$*, where* $\mathcal{A}_{i\mathcal{O},0}(1^\lambda)$ *outputs same-size circuits* $C_0,C_1\in\mathcal{C}_\lambda$ *and a quantum state* σ_{aux} *such that* $C_0(x)=C_1(x)$ *for every input* x*, there exists a negligible function* $\mathsf{negl}(\cdot)$ *such that*

$$\Pr\left[\widehat{b}=b:\begin{array}{l}(C_0,C_1,\sigma_{\mathsf{aux}})\leftarrow\mathcal{A}_{i\mathcal{O},0}(1^\lambda),\\ b\xleftarrow{\$}\{0,1\},\quad \rho_b\leftarrow\mathsf{Obf}(1^\lambda,C_b),\\ \widehat{b}\leftarrow\mathcal{A}_{i\mathcal{O},1}(\sigma_{\mathsf{aux}},\rho_b)\end{array}\right]\le\frac{1}{2}+\mathsf{negl}(\lambda). \tag{55}$$

Next, we show the following composition theorem for correlated security of UPO using the same argument as in the proof of [AB24, Theorem 12].

Theorem 18 (Adding $i\mathcal{O}$ security, adapted from [AB24, Theorem 12],[ABH+26, Lemma 3.16])**.** *Let* $\mathsf{Circ}_{\ell_{\mathsf{in}},\ell_{\mathsf{out}}}$ *denote the full class of polynomial-size circuits from* $\{0,1\}^{\ell_{\mathsf{in}}(\lambda)}$ *to* $\{0,1\}^{\ell_{\mathsf{out}}(\lambda)}$*, regarded as a keyed circuit class whose keys are circuit descriptions. If there exist a post-quantum secure indistinguishability obfuscation scheme and a UPO scheme for* $\mathsf{Circ}_{\ell_{\mathsf{in}},\ell_{\mathsf{out}}}$ *satisfying* κ*-entropic correlated UPO security, then there exists a UPO scheme for this circuit class satisfying both* κ*-entropic correlated UPO security and* $i\mathcal{O}$*-security. Consequently, the same holds for every polynomial-size keyed subclass* $\mathcal{C}\subseteq\mathsf{Circ}_{\ell_{\mathsf{in}},\ell_{\mathsf{out}}}$*.*

Proof. This follows from the add-$i\mathcal{O}$ transformation of [ABH+26, Lemma 3.16], using the composition theorem of [AB24, Theorem 12]; the same reduction applies to the correlated experiment by using the sampler's entire output tuple, including the pre-split and post-split auxiliary strings, unchanged. □

Lemma 6 (Relation to prior UPO notions)**.** *For any* $\kappa:=\kappa(\lambda)=\lambda^c$ *for some constant* $c>0$*,* κ*-entropic* Correlated-UPO *security implies* $\mathcal{D}_X$*-generalized UPO security for every QPT-samplable distribution* $\mathcal{D}_X$ *over challenge pairs, that induces marginal challenge distribution on each of the two challenges with min-entropy at least* $\kappa(\lambda)$*. In particular, for every QPT-samplable distribution* $\mathcal{D}$ *satisfying* $\widetilde{H}_\infty(\mathcal{D})\ge\kappa(\lambda)$ *(such as the uniform random distribution), it implies* $\mathcal{D}^2$*-generalized UPO security and* $\mathsf{ID}_{\mathcal{D}}$*-generalized UPO security.*[9]*.*

Proof. Given $\mathcal{D}_X$, define a sampler that sets $z_{\mathsf{aux}}:=w_{\mathcal{B}}:=w_{\mathcal{C}}:=\bot$, samples $(x_{\mathcal{B}},x_{\mathcal{C}})\leftarrow\mathcal{D}_X$, samples $b\xleftarrow{\$}\{0,1\}$, and sets $b_{\mathcal{B}}=b_{\mathcal{C}}:=b$. By the hypothesis of the lemma, $\mathcal{D}_X$, the resulting sampler has κ-entropic marginals on the individual challenges. Moreover, clearly, this sampler satisfies Equation (52). Since every nonempty sampler output is hidden from $\mathcal{A}_0$ before it sends $(k,\mu_{\mathcal{B}},\mu_{\mathcal{C}})$, sampling the tuple before this message does not change the joint distribution.

If $b=0$, the punctured circuit is C_k, while if $b=1$, it is exactly the generalized punctured circuit from Equation (50), including the Bob-priority convention used in the generalized UPO security definition when the points coincide. Hence the two security experiments are identically distributed. Taking $\mathcal{D}_X=\mathcal{D}^2$, $\mathcal{D}_X=\mathsf{ID}_{\mathcal{D}}$, $\mathcal{D}_X=\mathsf{U}$, or $\mathcal{D}_X=\mathsf{ID}_{\mathsf{U}}$ gives the stated instantiations. Vanilla UPO security follows since it is implied by generalized UPO security. □

[9] The only known UPO security notion that is not covered by correlated UPO security is the UPO+ notion defined in [ABH+26], which can be seen as a completely different flavor of security notion. In particular, UPO$^+$ has a different winning condition and is not an instance of the correlated-UPO experiment.

5.2 Copy-Protection Security Definitions

We first recall the standard copy-protection security notions, including the identical-challenge point-function experiment of [AB24]. We then define the stronger oracular variants of [ABH+26]. Except in the point-function, k-point-function, and compute-and-compare experiments stated separately below, the circuit and the challenge sampler are independent, while the two challenge points themselves may be arbitrarily correlated.

Let $\mathsf{Circ} = \{\mathsf{Circ}_\lambda\}_{\lambda\in\mathbb{N}}$ be a class of polynomial-size circuits $C : \mathcal{X}_\lambda \to \mathcal{Y}_\lambda$, where $\mathcal{X}_\lambda = \{0,1\}^{\ell_{\mathsf{in}}(\lambda)}$ and $\mathcal{Y}_\lambda = \{0,1\}^{\ell_{\mathsf{out}}(\lambda)}$. Let $\mathcal{D}_{\mathsf{Circ}}$ be a QPT-samplable distribution over Circ_λ.

Definition 29 (Copy protection, adapted from [AB24, Appendix A.1] and [ABH+26, Definition 3.1])**.** *A copy-protection scheme* $\mathsf{CP} = (\mathsf{CopyProtect}, \mathsf{Eval})$ *for* Circ *consists of a pair of efficient algorithms as follows.*

1. $\rho_C \leftarrow \mathsf{CopyProtect}(1^\lambda, C)$: *is a QPT algorithm that takes as input a circuit* $C \in \mathsf{Circ}_\lambda$ *and outputs a quantum program* ρ_C*.*
2. $y \leftarrow \mathsf{Eval}(\rho_C, x)$: *is a QPT algorithm that takes as input a quantum program* ρ_C *and an input* $x \in \mathcal{X}_\lambda$*, and outputs* $y \in \mathcal{Y}_\lambda$*.*

Correctness. *We require that there exists a negligible function* $\mathsf{negl}(\cdot)$ *such that, for every* $\lambda \in \mathbb{N}$*, every* $C \in \mathsf{Circ}_\lambda$*, and every* $x \in \mathcal{X}_\lambda$*,*

$$\Pr\left[\mathsf{Eval}(\rho_C, x) = C(x) : \rho_C \leftarrow \mathsf{CopyProtect}(1^\lambda, C)\right] \geq 1 - \mathsf{negl}(\lambda). \tag{56}$$

When polynomial reusability is required, the evaluation algorithm is understood to return its residual program state; the concrete scheme used below restores that state perfectly.

Definition 30 (Correlated copy-protection challenge sampler)**.** *A correlated copy-protection challenge sampler* $\mathsf{Samp}_{\mathsf{CP}}$ *is a QPT algorithm that, on input* 1^λ*, outputs* $(x_{\mathcal{B}}, x_{\mathcal{C}}) \leftarrow \mathsf{Samp}_{\mathsf{CP}}(1^\lambda)$*, where* $x_{\mathcal{B}}, x_{\mathcal{C}} \in \mathcal{X}_\lambda$*. The two challenge points may be arbitrarily correlated. We say that* $\mathsf{Samp}_{\mathsf{CP}}$ *has* κ*-entropic marginals if* $\widetilde{H}_\infty(x_X) \geq \kappa(\lambda)$ *for every* $X \in \{\mathcal{B}, \mathcal{C}\}$*.*

We use the notion of trivial adversaries from previous works [ABH+26, Definition 3.1]. The first-stage adversary $\mathcal{A}^{\mathsf{triv}}_{\mathcal{B}}$ sends the entire protected program ρ_C to $\mathcal{B}$ and the fixed dummy state $|\bot\rangle\langle\bot|$ to $\mathcal{C}$. The first-stage adversary $\mathcal{A}^{\mathsf{triv}}_{\mathcal{C}}$ is defined symmetrically. Define

$$\mathsf{Triv} := \left\{(\mathcal{A}^{\mathsf{triv}}_{\mathcal{B}}, \mathcal{B}, \mathcal{C}) : \mathcal{B}, \mathcal{C} \text{ are QPT}\right\} \cup \left\{(\mathcal{A}^{\mathsf{triv}}_{\mathcal{C}}, \mathcal{B}, \mathcal{C}) : \mathcal{B}, \mathcal{C} \text{ are QPT}\right\}. \tag{57}$$

Thus, the recipient not receiving the protected program receives no auxiliary state prepared by $\mathcal{A}$.

For any QPT adversarial triplet $\mathcal{A}_{\mathsf{CP}} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$, consider the following experiment.
$\mathsf{UnpCP}_{\mathsf{CP},\mathcal{A}_{\mathsf{CP}},\mathcal{D}_{\mathsf{Circ}},\mathsf{Samp}_{\mathsf{CP}}}(1^\lambda)$:

1. Ch samples $C \leftarrow \mathcal{D}_{\mathsf{Circ}}(1^\lambda)$ and $\rho_C \leftarrow \mathsf{CopyProtect}(1^\lambda, C)$, and gives ρ_C to $\mathcal{A}$.
2. $\mathcal{A}$ outputs a joint state $\sigma_{\mathsf{R}_1,\mathsf{R}_2}$ and sends register R_1 to $\mathcal{B}$ and register R_2 to $\mathcal{C}$.
3. After the split, Ch samples $(x_{\mathcal{B}}, x_{\mathcal{C}}) \leftarrow \mathsf{Samp}_{\mathsf{CP}}(1^\lambda)$, gives $x_{\mathcal{B}}$ to $\mathcal{B}$, and gives $x_{\mathcal{C}}$ to $\mathcal{C}$.
4. $\mathcal{B}$ and $\mathcal{C}$ output $\widehat{y}_{\mathcal{B}}, \widehat{y}_{\mathcal{C}}$.
5. The output of the experiment is 1 if $\widehat{y}_{\mathcal{B}} = C(x_{\mathcal{B}})$ and $\widehat{y}_{\mathcal{C}} = C(x_{\mathcal{C}})$.

Definition 31 (Unpredictability-style copy-protection security with correlated challenges, adapted from [AB24, Appendix A.1] and [ABH+26, Definition 3.1])**.** *A copy-protection scheme* CP *satisfies unpredictability-style copy-protection security with respect to* $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{CP}})$ *if, for every QPT adversarial triplet* $\mathcal{A}_{\mathsf{CP}}$*, there exist a trivial adversary* $\mathcal{A}_{\mathsf{triv}} \in \mathsf{Triv}$ *and a negligible function* $\mathsf{negl}(\cdot)$ *such that*

$$\Pr\left[\mathsf{UnpCP}_{\mathsf{CP},\mathcal{A}_{\mathsf{CP}},\mathcal{D}_{\mathsf{Circ}},\mathsf{Samp}_{\mathsf{CP}}}(1^\lambda) = 1\right] \leq \Pr\left[\mathsf{UnpCP}_{\mathsf{CP},\mathcal{A}_{\mathsf{triv}},\mathcal{D}_{\mathsf{Circ}},\mathsf{Samp}_{\mathsf{CP}}}(1^\lambda) = 1\right] + \mathsf{negl}(\lambda). \tag{58}$$

Definition 32 (Identical-challenge unpredictability-style copy protection for point functions, adapted from [AB24, Theorem 89 and Corollary 90])**.** *Let* $\mathcal{X}_\lambda = \{0,1\}^{\ell_{\mathsf{in}}(\lambda)}$*, let* $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda\in\mathbb{N}}$ *be a QPT-samplable distribution over* $\mathcal{X}_\lambda$*, and for every* $y \in \mathcal{X}_\lambda$ *define* $P_y(x) := [x = y]$*. For a copy-protection scheme* CP *for point functions and any QPT adversarial triplet* $\mathcal{A}_{\mathsf{PF}} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$*, consider the following experiment.*
$\mathsf{PFCP}_{\mathsf{CP},\mathcal{A}_{\mathsf{PF}},\mathcal{D}}(1^\lambda)$*:*

1. *Ch samples* $y \leftarrow \mathcal{D}_\lambda$ *and* $\rho_y \leftarrow \mathsf{CopyProtect}(1^\lambda, P_y)$*, and sends* ρ_y *to* $\mathcal{A}$*.*
2. $\mathcal{A}$ *outputs a joint state* $\sigma_{\mathsf{R}_1,\mathsf{R}_2}$ *and sends register* R_1 *to* $\mathcal{B}$ *and register* R_2 *to* $\mathcal{C}$*.*
3. *After the split, Ch samples* $b \xleftarrow{\$} \{0,1\}$*. If* $b = 0$*, it samples an independent* $x \leftarrow \mathcal{D}_\lambda$*; if* $b = 1$*, it sets* $x := y$*. It sends the same* x *to both* $\mathcal{B}$ *and* $\mathcal{C}$*.*
4. $\mathcal{B}$ *and* $\mathcal{C}$ *output* $\widehat{v}_{\mathcal{B}}, \widehat{v}_{\mathcal{C}} \in \{0,1\}$*.*
5. *The output of the experiment is* 1 *if* $\widehat{v}_{\mathcal{B}} = \widehat{v}_{\mathcal{C}} = P_y(x)$*.*

A copy-protection scheme for point functions satisfies identical-challenge point-function security with respect to $\mathcal{D}$ *if, for every QPT adversarial triplet* $\mathcal{A}_{\mathsf{PF}}$*, there exist a trivial adversary* $\mathcal{A}_{\mathsf{triv}} \in \mathsf{Triv}$ *and a negligible function* $\mathsf{negl}(\cdot)$ *such that*

$$\Pr\left[\mathsf{PFCP}_{\mathsf{CP},\mathcal{A}_{\mathsf{PF}},\mathcal{D}}(1^\lambda) = 1\right] \leq p_{\mathsf{triv}}^{\mathsf{PF},\mathcal{D}}(\lambda) + \mathsf{negl}(\lambda), \tag{59}$$

where

$$p_{\mathsf{triv}}^{\mathsf{PF},\mathcal{D}}(\lambda) := \sup_{\mathcal{A}_{\mathsf{triv}}\in\mathsf{Triv}} \Pr\left[\mathsf{PFCP}_{\mathsf{CP},\mathcal{A}_{\mathsf{triv}},\mathcal{D}}(1^\lambda) = 1\right]. \tag{60}$$

The uniform choice of $\mathcal{D}_\lambda$ *recovers the experiment of [AB24, Theorem 89 and Corollary 90].*

We note that, assuming the point distribution has sufficiently high min-entropy, the trivial success probability is negligibly close to $1/2$.

Lemma 7 (Trivial success in the point-function experiment)**.** *Let* $\mathrm{cp}(\mathcal{D}_\lambda) := \sum_x \Pr[\mathcal{D}_\lambda = x]^2$*. For the experiment in Definition 32,*

$$p_{\mathsf{triv}}^{\mathsf{PF},\mathcal{D}}(\lambda) = \frac{1}{2} + \frac{1}{2}\,\mathrm{cp}(\mathcal{D}_\lambda) \leq \frac{1}{2} + \frac{1}{2}\,2^{-\widetilde{H}_\infty(\mathcal{D}_\lambda)}. \tag{61}$$

Consequently, if $\widetilde{H}_\infty(\mathcal{D}_\lambda) \geq \lambda^c$ *for some constant* $c > 0$*, then*

$$p_{\mathsf{triv}}^{\mathsf{PF},\mathcal{D}}(\lambda) = \frac{1}{2} + \mathsf{negl}'(\lambda). \tag{62}$$

Proof. Both events induce the marginal distribution $x \leftarrow \mathcal{D}_\lambda$. For every x in its support,

$$\Pr[P_y(x) = 1 \mid x] = \frac{1}{2}\big(1 + \Pr[\mathcal{D}_\lambda = x]\big). \tag{63}$$

Under a trivial strategy, the recipient receiving the dummy register has no information about y beyond x, and hence cannot do better than always outputting 1. This strategy is attained when both recipients output 1 and has average success $1/2 + \mathrm{cp}(\mathcal{D}_\lambda)/2$. Finally, $\mathrm{cp}(\mathcal{D}_\lambda) \le \max_x \Pr[\mathcal{D}_\lambda = x] = 2^{-\widetilde{H}_\infty(\mathcal{D}_\lambda)}$. □

Remark 4 (Copy protection definition for k-point functions)**.** *The copy protection definition for point functions, as outlined above, also extends to the fixed-k-point functions considered in [AB24, Theorem 85 and Corollary 90]. Let $\mathcal{D}^{(k)} = \{\mathcal{D}_\lambda^{(k)}\}_{\lambda\in\mathbb{N}}$ be any QPT-samplable distribution over the $k = k(\lambda)$ element subsets of $\mathcal{X}_\lambda$ such that $\widetilde{H}_\infty(\mathcal{D}_\lambda^{(k)}) \ge \lambda^c$ for some constant $c > 0$. For a set $S \subseteq \mathcal{X}_\lambda$ of size k, define $P_S(x) := [x \in S]$, and let $\mathcal{D}_\lambda^{\mathsf{pt}}$ denote the distribution obtained by sampling $S \leftarrow \mathcal{D}_\lambda^{(k)}$ and then $x \xleftarrow{\$} S$, such that $\widetilde{H}_\infty(\mathcal{D}_\lambda^{\mathsf{pt}}) \ge \lambda^{c'}$ for some constant $c' > 0$. The identical-challenge experiment samples $S \leftarrow \mathcal{D}_\lambda^{(k)}$ and protects P_S. After the split, it samples $b \xleftarrow{\$} \{0,1\}$. If $b = 0$, it samples an independent $S' \leftarrow \mathcal{D}_\lambda^{(k)}$ and then $x \xleftarrow{\$} S'$; if $b = 1$, it samples $x \xleftarrow{\$} S$. In either case, it gives the same x to both recipients. Security uses the same trivial-adversary comparison as above.*

For $q_\lambda(x) := \Pr_{S\leftarrow\mathcal{D}_\lambda^{(k)}}[x \in S]$, the trivial success probability is

$$p_{\mathsf{triv}}^{\mathsf{PF}(k),\mathcal{D}^{(k)}}(\lambda) = \frac{1}{2} + \frac{1}{2k(\lambda)} \sum_{x\in\mathcal{X}_\lambda} q_\lambda(x)^2 = \frac{1}{2} + \frac{k(\lambda)}{2}\,\mathrm{cp}(\mathcal{D}_\lambda^{\mathsf{pt}}), \tag{64}$$

which is $1/2 + \mathsf{negl}(\lambda)$ whenever k is polynomially bounded.

Definition 33 (Identical-challenge copy protection for compute-and-compare functions, adapted from [CMP24])**.** *Following [CMP24, Construction 4 and Theorem 7], the definition above extends to compute-and-compare functions*

$$\mathsf{CC}_{f,y}(x) := [f(x) = y],$$

where $f : \mathcal{X}_\lambda \to \mathcal{Z}_\lambda$, $\mathcal{Z}_\lambda = \{0,1\}^{\ell_{\mathsf{cc}}(\lambda)}$, and $y \in \mathcal{Z}_\lambda$. Let Samp be a QPT algorithm that samples $(f, y) \leftarrow \mathsf{Samp}(1^\lambda)$, interpreted as sampling $\mathsf{CC}_{f,y}$, and let (F, Y) denote its output distribution. We require $\widetilde{H}_\infty(Y \mid F) \ge \lambda^c$ for some constant $c > 0$.

For every f in the support of F and every z in the support of the conditional distribution of Y given $F = f$, fix a distribution $\mathcal{Q}_{f,z}$ supported on $f^{-1}(z)$. The identical-challenge copy-protection experiment for compute-and-compare functions samples $(f, y) \leftarrow \mathsf{Samp}(1^\lambda)$ and protects $\mathsf{CC}_{f,y}$ as (f, ρ_y), where ρ_y is a copy protection of P_y. After the split, it samples $b \xleftarrow{\$} \{0,1\}$. If $b = 0$, it samples a fresh z independently according to the conditional distribution of Y given $F = f$; if $b = 1$, it sets $z := y$. It then samples $x \leftarrow \mathcal{Q}_{f,z}$ and gives the same x to both recipients. The same family of distributions $\{\mathcal{Q}_{f,z}\}_{f,z}$ is used for both values of b. We say that the scheme satisfies identical-challenge compute-and-compare copy-protection security if, for every QPT adversarial triplet, there exists a negligible function negl such that its success probability in the experiment above is at most

$$p_{\mathsf{triv}}^{\mathsf{CC},\mathsf{Samp}}(\lambda) + \mathsf{negl}(\lambda),$$

where $p_{\mathsf{triv}}^{\mathsf{CC},\mathsf{Samp}}(\lambda)$ *denotes the supremum success probability in this experiment over the trivial adversaries in* Triv*, as defined in Definition 32.*

We say that this identical-challenge experiment is efficiently realizable if there exist QPT procedures that:

1. *on input* $(1^\lambda, f)$*, sample a fresh* z *according to the conditional distribution of* Y *given* $F = f$*; and*
2. *on input* $(1^\lambda, f, z)$*, sample* $x \leftarrow \mathcal{Q}_{f,z}$*.*

Next, we recall the pseudorandomness-style copy protection notion from previous works.

Definition 34 (Correlated pseudorandomness-style copy-protection challenge sampler)**.** *A correlated pseudorandomness-style copy-protection challenge sampler* $\mathsf{Samp}_{\mathsf{PRCP}}$ *with* κ*-entropic marginals is a QPT algorithm that, on input* 1^λ*, outputs* $(x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}) \leftarrow \mathsf{Samp}_{\mathsf{PRCP}}(1^\lambda)$*, where* $x_{\mathcal{B}}, x_{\mathcal{C}} \in \mathcal{X}_\lambda$ *and* $b_{\mathcal{B}}, b_{\mathcal{C}} \in \{0,1\}$*, such that, for every* $(\bar{x}_{\mathcal{B}}, \bar{x}_{\mathcal{C}})$ *in the support of* $(x_{\mathcal{B}}, x_{\mathcal{C}})$ *and every* $X \in \{\mathcal{B}, \mathcal{C}\}$*,*

$$\Pr[b_X = 0 \mid x_{\mathcal{B}} = \bar{x}_{\mathcal{B}}, x_{\mathcal{C}} = \bar{x}_{\mathcal{C}}] = \Pr[b_X = 1 \mid x_{\mathcal{B}} = \bar{x}_{\mathcal{B}}, x_{\mathcal{C}} = \bar{x}_{\mathcal{C}}] = \frac{1}{2}, \tag{65}$$

and $\widetilde{H}_\infty(x_X) \geq \kappa(\lambda)$ *for every* $X \in \{\mathcal{B}, \mathcal{C}\}$*.*

We write $\mathsf{Samp}_{\mathsf{ID_U}}^{\mathsf{pr}}$ for the sampler that samples $x \xleftarrow{\$} \mathcal{X}_\lambda$ and $b \xleftarrow{\$} \{0,1\}$, and outputs $(x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}) := (x, x, b, b)$.

Given C, a sampled tuple $(x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}})$, and independent $y_{\mathcal{B}}, y_{\mathcal{C}} \xleftarrow{\$} \mathcal{Y}_\lambda$, define

$$z_{\mathcal{B}} := \begin{cases} C(x_{\mathcal{B}}), & b_{\mathcal{B}} = 0, \\ y_{\mathcal{B}}, & b_{\mathcal{B}} = 1, \end{cases} \qquad z_{\mathcal{C}} := \begin{cases} z_{\mathcal{B}}, & x_{\mathcal{C}} = x_{\mathcal{B}}, \\ C(x_{\mathcal{C}}), & x_{\mathcal{C}} \neq x_{\mathcal{B}} \text{ and } b_{\mathcal{C}} = 0, \\ y_{\mathcal{C}}, & x_{\mathcal{C}} \neq x_{\mathcal{B}} \text{ and } b_{\mathcal{C}} = 1. \end{cases} \tag{66}$$

Thus, just as in the correlated UPO security case (see Equation (53)), a potential tie (i.e., $x_{\mathcal{B}} = x_{\mathcal{C}}$ but $b_{\mathcal{B}} \neq b_{\mathcal{C}}$) is resolved using $\mathcal{B}$'s challenge bit.

For any QPT adversarial triplet $\mathcal{A}_{\mathsf{PRCP}} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$, consider the following experiment. $\mathsf{PRCP}_{\mathsf{CP},\mathcal{A}_{\mathsf{PRCP}},\mathcal{D}_{\mathsf{Circ}},\mathsf{Samp}_{\mathsf{PRCP}}}(1^\lambda)$:

1. Ch samples $C \leftarrow \mathcal{D}_{\mathsf{Circ}}(1^\lambda)$ and $\rho_C \leftarrow \mathsf{CopyProtect}(1^\lambda, C)$, and gives ρ_C to $\mathcal{A}$.
2. $\mathcal{A}$ outputs a joint state $\sigma_{\mathsf{R}_1,\mathsf{R}_2}$ and sends register R_1 to $\mathcal{B}$ and register R_2 to $\mathcal{C}$.
3. After the split, Ch samples $(x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}) \leftarrow \mathsf{Samp}_{\mathsf{PRCP}}(1^\lambda)$ and independent $y_{\mathcal{B}}, y_{\mathcal{C}} \xleftarrow{\$} \mathcal{Y}_\lambda$. It defines $z_{\mathcal{B}}, z_{\mathcal{C}}$ according to Equation (66).
4. Ch gives $(x_{\mathcal{B}}, z_{\mathcal{B}})$ to $\mathcal{B}$ and $(x_{\mathcal{C}}, z_{\mathcal{C}})$ to $\mathcal{C}$.
5. $\mathcal{B}$ and $\mathcal{C}$ output $\widehat{b}_{\mathcal{B}}, \widehat{b}_{\mathcal{C}}$.
6. The output of the experiment is 1 if $\widehat{b}_{\mathcal{B}} = b_{\mathcal{B}}$ and $\widehat{b}_{\mathcal{C}} = b_{\mathcal{C}}$.

Definition 35 (Pseudorandomness-style copy-protection security with correlated challenges, adapted from [ABH+26, Definition 3.3])**.** *A copy-protection scheme* CP *satisfies* κ*-entropic pseudorandomness-style copy-protection security with respect to* $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{PRCP}})$ *if, for every QPT adversarial triplet* $\mathcal{A}_{\mathsf{PRCP}}$*, there exists a negligible function* $\mathsf{negl}(\cdot)$ *such that*

$$\Pr\left[\mathsf{PRCP}_{\mathsf{CP},\mathcal{A}_{\mathsf{PRCP}},\mathcal{D}_{\mathsf{Circ}},\mathsf{Samp}_{\mathsf{PRCP}}}(1^\lambda) = 1\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda). \tag{67}$$

5.2.1 Oracular security notions

Next we recall the oracular versions of the copy-protection security definitions defined in [ABH+26]. The experiments below are identical to the experiments above, except that $\mathcal{A}$ receives oracle access to the sampled circuit before the split and each recipient receives oracle access punctured at that recipient's challenge point after the split.

For every circuit $C : \mathcal{X}_\lambda \to \mathcal{Y}_\lambda$ and $x^\star \in \mathcal{X}_\lambda$, define

$$C^{\setminus x^\star}(x) := \begin{cases} \bot, & x = x^\star, \\ C(x), & x \neq x^\star. \end{cases} \tag{68}$$

All oracle accesses below are quantum superposition accesses.

For any QPT adversarial triplet $\mathcal{A}_{\mathsf{CP}} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$, consider the following experiment. $\mathsf{OraUnpCP}_{\mathsf{CP},\mathcal{A}_{\mathsf{CP}},\mathcal{D}_{\mathsf{Circ}},\mathsf{Samp}_{\mathsf{CP}}}(1^\lambda)$:

1. Ch samples $C \leftarrow \mathcal{D}_{\mathsf{Circ}}(1^\lambda)$ and $\rho_C \leftarrow \mathsf{CopyProtect}(1^\lambda, C)$, and gives ρ_C and oracle access to C to $\mathcal{A}$.
2. $\mathcal{A}^C$ outputs a joint state $\sigma_{\mathsf{R}_1,\mathsf{R}_2}$ and sends register R_1 to $\mathcal{B}$ and register R_2 to $\mathcal{C}$.
3. After the split, Ch samples $(x_\mathcal{B}, x_\mathcal{C}) \leftarrow \mathsf{Samp}_{\mathsf{CP}}(1^\lambda)$, gives $x_\mathcal{B}$ and oracle access to $C^{\setminus x_\mathcal{B}}$ to $\mathcal{B}$, and gives $x_\mathcal{C}$ and oracle access to $C^{\setminus x_\mathcal{C}}$ to $\mathcal{C}$.
4. $\mathcal{B}^{C^{\setminus x_\mathcal{B}}}$ and $\mathcal{C}^{C^{\setminus x_\mathcal{C}}}$ output $\widehat{y}_\mathcal{B}, \widehat{y}_\mathcal{C}$.
5. The output of the experiment is 1 if $\widehat{y}_\mathcal{B} = C(x_\mathcal{B})$ and $\widehat{y}_\mathcal{C} = C(x_\mathcal{C})$.

The notion of trivial strategies in this experiment is defined similarly to the case of no oracle access. For $X \in \{\mathcal{B}, \mathcal{C}\}$, rhe first-stage adversary $\mathcal{A}^{\mathsf{triv}}_X$ with access to the oracle sends the entire protected program ρ_C to X and the fixed dummy state $|\bot\rangle\langle\bot|$ to the other party, and, we define

$$\mathsf{Triv} := \left\{(\mathcal{A}^{\mathsf{triv}}_\mathcal{B}, \mathcal{B}, \mathcal{C}) : \mathcal{B}, \mathcal{C} \text{ are QPT}\right\} \cup \left\{(\mathcal{A}^{\mathsf{triv}}_\mathcal{C}, \mathcal{B}, \mathcal{C}) : \mathcal{B}, \mathcal{C} \text{ are QPT}\right\}. \tag{69}$$

Definition 36 (Oracular unpredictability-style copy-protection security with correlated challenges, adapted from [ABH+26, Definition 3.1])**.** *A copy-protection scheme* CP *satisfies oracular unpredictability-style copy-protection security with respect to* $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{CP}})$ *if, for every QPT adversarial triplet* $\mathcal{A}_{\mathsf{CP}}$*, there exist a trivial adversary* $\mathcal{A}_{\mathsf{triv}} \in \mathsf{Triv}$ *(defined in Equation (69)) and a negligible function* $\mathsf{negl}(\cdot)$ *such that*

$$\Pr\left[\mathsf{OraUnpCP}_{\mathsf{CP},\mathcal{A}_{\mathsf{CP}},\mathcal{D}_{\mathsf{Circ}},\mathsf{Samp}_{\mathsf{CP}}}(1^\lambda) = 1\right] \leq \Pr\left[\mathsf{OraUnpCP}_{\mathsf{CP},\mathcal{A}_{\mathsf{triv}},\mathcal{D}_{\mathsf{Circ}},\mathsf{Samp}_{\mathsf{CP}}}(1^\lambda) = 1\right] + \mathsf{negl}(\lambda). \tag{70}$$

For any QPT adversarial triplet $\mathcal{A}_{\mathsf{PRCP}} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$, consider the following experiment. $\mathsf{OraPRCP}_{\mathsf{CP},\mathcal{A}_{\mathsf{PRCP}},\mathcal{D}_{\mathsf{Circ}},\mathsf{Samp}_{\mathsf{PRCP}}}(1^\lambda)$:

1. Ch samples $C \leftarrow \mathcal{D}_{\mathsf{Circ}}(1^\lambda)$ and $\rho_C \leftarrow \mathsf{CopyProtect}(1^\lambda, C)$, and gives ρ_C and oracle access to C to $\mathcal{A}$.
2. $\mathcal{A}^C$ outputs a joint state $\sigma_{\mathsf{R}_1,\mathsf{R}_2}$ and sends register R_1 to $\mathcal{B}$ and register R_2 to $\mathcal{C}$.
3. After the split, Ch samples $(x_\mathcal{B}, x_\mathcal{C}, b_\mathcal{B}, b_\mathcal{C}) \leftarrow \mathsf{Samp}_{\mathsf{PRCP}}(1^\lambda)$ and independent $y_\mathcal{B}, y_\mathcal{C} \xleftarrow{\$} \mathcal{Y}_\lambda$. It defines $z_\mathcal{B}, z_\mathcal{C}$ according to Equation (66).

4. Ch gives $(x_{\mathcal{B}}, z_{\mathcal{B}})$ and oracle access to $C^{\setminus x_{\mathcal{B}}}$ to $\mathcal{B}$, and gives $(x_{\mathcal{C}}, z_{\mathcal{C}})$ and oracle access to $C^{\setminus x_{\mathcal{C}}}$ to $\mathcal{C}$.

5. $\mathcal{B}^{C^{\setminus x_{\mathcal{B}}}}$ and $\mathcal{C}^{C^{\setminus x_{\mathcal{C}}}}$ output $\widehat{b}_{\mathcal{B}}, \widehat{b}_{\mathcal{C}}$.

6. The output of the experiment is 1 if $\widehat{b}_{\mathcal{B}} = b_{\mathcal{B}}$ and $\widehat{b}_{\mathcal{C}} = b_{\mathcal{C}}$.

Definition 37 (Oracular pseudorandomness-style copy-protection security with correlated challenges, adapted from [ABH+26, Definition 3.3])**.** *A copy-protection scheme* CP *satisfies* κ*-entropic oracular pseudorandomness-style copy-protection security with respect to* $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{PRCP}})$ *if, for every QPT adversarial triplet* $\mathcal{A}_{\mathsf{PRCP}}$*, there exists a negligible function* $\mathsf{negl}(\cdot)$ *such that*

$$\Pr\left[\mathsf{OraPRCP}_{\mathsf{CP},\mathcal{A}_{\mathsf{PRCP}},\mathcal{D}_{\mathsf{Circ}},\mathsf{Samp}_{\mathsf{PRCP}}}(1^\lambda) = 1\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda). \tag{71}$$

Part II
Quantum Information Tools

6 A New Decisional Coset Monogamy-of-Entanglement Game

In this section, we state and prove a new information theoretic decisional Coset Monogamy-of-Entanglement game, that implies (with other cryptographic assumptions) all the feasibility results in this article, i.e., all the security proofs finally reduce to this Monogamy-of-Entanglement game. We note that all adversaries in this section are computationally unbounded, and can be modeled as arbitrary (finite-dimensional) quantum channels, and all the results hold unconditionally, without any computational assumptions.

Definition 38 (Decisional-Coset-MOE)**.** *Let $n := n(\lambda)$ be a polynomial function in λ. For any triplet of adversaries $\mathcal{A}_{moe} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$, consider the following experiment between a challenger Ch and $\mathcal{A}_{moe}$.*
Decisional-Coset-MOE$_{\mathcal{A}_{moe}}(1^\lambda)$:

1. *Ch samples a uniformly random n-dimensional subspace $A \leq \mathbb{F}_2^{2n}$ and samples $s, t \xleftarrow{\$} \mathbb{F}_2^{2n}$ independently, and prepares and sends the state $|A_{s,t}\rangle = \frac{1}{\sqrt{|A|}} \sum_{a \in A} (-1)^{\langle a,t \rangle} |a + s\rangle$.*
2. *$\mathcal{A}$ outputs a joint state $\sigma_{\mathsf{R}_1,\mathsf{R}_2}$ and sends registers R_1 to $\mathcal{B}$ and R_2 to $\mathcal{C}$.*
3. *After the split, Ch samples $u \xleftarrow{\$} A \setminus \{0\}$ and $v \xleftarrow{\$} A^\perp \setminus \{0\}$ and sends the pair (u, v) to $\mathcal{B}$ and $\mathcal{C}$.*
4. *$\mathcal{B}$ and $\mathcal{C}$ output $b_\mathcal{B}, b_\mathcal{C}$.*
5. *The output of the experiment 1 if $b_\mathcal{B} = b_\mathcal{C} = \langle u, t \rangle + \langle v, s \rangle$.*

It is easy to see that the Decisional-Coset-MOE experiment can be seen as an unclonable-indistinguishable experiment (see Definition 10) as follows.

Lemma 8. *The Decisional-Coset-MOE experiment is equivalent to the unclonable-indistinguishable experiment (see Definition 10) for the following unclonable encryption scheme.*
Coset-UE:

1. $\mathsf{GenKey}(1^\lambda)$*: Sample $u, v \leftarrow \mathbb{F}_2^{2n} \setminus \{0\}$ subject to $\langle u, v \rangle = 0$. Output $k = (u, v)$.*
2. $\mathsf{Enc}(k, m)$*: Sample a uniformly random subspace $A \leq \mathbb{F}_2^{2n}$ subject to $u \in A$ and $v \in A^\perp$ where $k = (u, v)$. Sample $s, t \xleftarrow{\$} \mathbb{F}_2^{2n}$ subject to $m = \langle u, t \rangle + \langle v, s \rangle$. Prepare and output the state $\rho_{ct} = |A_{s,t}\rangle$.*
3. $\mathsf{Dec}(k, \rho_{ct})$*: Interpret $k := (u, v)$, compute $m_1 = \langle v, s \rangle$, by computing the inner product of v coherently with all strings in superposition in the state $|A_{s,t}\rangle$ and then measure the outcome, which does not collapse the state (since $\langle a + s, v = \langle s, v \rangle$ for every $a \in A$). Then apply the QFT on the state to get the state $\left|A^\perp_{t,s}\right\rangle$, and then compute $m_2 = \langle u, t \rangle$ similarly using the state $\left|A^\perp_{t,s}\right\rangle$ and u. Output $m = m_1 + m_2$.*

Proof. For every fixed nonzero orthogonal pair $(u, v) \in \mathbb{F}_2^{2n} \times \mathbb{F}_2^{2n}$, let the set of n-dimensional subspaces A satisfying $u \in A$ and $v \in A^\perp$ be defined as Compatible(u, v). It is easy to see that |Compatible(u, v)| depends only on n and is independent of the choice of (u, v). Hence the marginal distribution on the pair (u, v) induced by the Decisional-Coset-MOE experiment is uniform over the key space of Coset-UE, given as $\mathcal{K} := \{(u, v) \mid u, v \in \mathbb{F}_2^{2n} \setminus \{0\}, \langle u, v\rangle = 0\}$. Thus sampling A uniformly at random, and then sampling $u \xleftarrow{\$} A \setminus \{0\}, v \xleftarrow{\$} A^\perp \setminus \{0\}$, is the same as first sampling a uniformly random nonzero orthogonal pair $(u, v) \in \mathbb{F}_2^{2n} \times \mathbb{F}_2^{2n}$, and then sampling $A \in$ Compatible(u, v), and therefore the marginal distribution of $(A, u, v))$ in the Decisional-Coset-MOE experiment and the unclonable-indistinguishable experiment (see Definition 10) for Coset-UE, is the same. Next, for a fixed non-zero orthogonal pair (u, v) and a fixed subspace $A \in$ Compatible(u, v), the map

$$(s, t) \longmapsto \langle u, t\rangle + \langle v, s\rangle$$

is a nonzero linear function and hence is balanced; therefore, sampling s, t uniformly and setting the message to this parity is equivalent to first sampling a uniform message m and then sampling s, t uniformly conditioned on

$$m = \langle u, t\rangle + \langle v, s\rangle.$$

Therefore, conditioned on $A, (u, v)$, the distribution of the subspace s, t and the bit $\langle u, t\rangle + \langle v, s\rangle$ in the Decisional-Coset-MOE experiment, is the same as that of s, t and the message bit m in the unclonable indistinguishable experiment for Coset-UE, which concludes the proof of the lemma. □

Theorem 19 (Decisional Coset Monogamy)**.** *For any $\lambda \in \mathbb{N}$, and any triplet of adversary $\mathcal{A}_{moe} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$,*

$$\Pr\left[\textit{Decisional-Coset-MOE}_{\mathcal{A}_{moe}}(1^\lambda) = 1\right] \leq \frac{1}{2} + \frac{1}{2}\sqrt{\frac{2^{2n}}{(2^{2n}-1)(2^{2n-1}-1)}}.$$

Due to Lemma 8, Theorem 19 gives the following immediate corollary.

Corollary 1. *The encryption scheme Coset-UE is an unconditional $\frac{1}{2}\sqrt{\frac{2^{2n}}{(2^{2n}-1)(2^{2n-1}-1)}}$-unclonable indistinguishable encryption scheme.*

In order to prove Theorem 19, we will use its interpretation as a unclonable indistinguishable experiment, given in Lemma 8. In particular, we mildly extend the proof of [AS26][10] to get a general lemma showing unclonable-indistinguishable security of a class of unclonable bit-encryption schemes, that captures both Coset-UE and the construction of [AS26].

Lemma 9 (General class of unclonable indistinguishable encryption schemes)**.** *Suppose for every $\ell \in \mathbb{N}$, there exists a set of keys $\mathcal{K}$ and a set of hermitian operators $\{\Theta_k\}_{k\in\mathcal{K}}$ in $\mathcal{L}((\mathbb{C}^2)^{\otimes\ell})$, indexed by $\mathcal{K}$, satisfying that for every k, $\mathrm{tr}[\Theta_k] = 0$, $\Theta_k^2 = I$, and for every $k, k' \in \mathcal{K}$, $\mathrm{tr}[\Theta_k\Theta_{k'}] = N\delta_{k,k'}$ where $N = 2^\ell$ is the dimension of the space, and δ_{ij} is the Kronecker Delta function, then for any distribution $P = \{p_k\}_{k\in\mathcal{K}}$ on $\mathcal{K}$, the following quantum encryption scheme given by*

1. $\mathsf{GenKey}(1^\lambda)$*: Sample $k \leftarrow P$, and output k.*
2. $\mathsf{Enc}(k, m)$*: Prepare and output the state $\rho_{ct}^{k,m} = 1/N\,(I + (-1)^m\Theta_k)$.*
3. $\mathsf{Dec}(k, \rho)$*: Measure ρ in the eigen basis of Θ_k, i.e., the POVM $(\Pi_0, \Pi_1) := \left(\frac{I+\Theta_k}{2}, \frac{I-\Theta_k}{2}\right)$ and output the outcome.*

[10] A similar generalization of [AS26] have been independently shown in the concurrent works of [BBC26, BDF26]

satisfies that for every λ, and adversarial triplet $\mathcal{A}_{\mathsf{UE}} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ against it in the unclonable-indistinguishable experiment (see Definition 10), the adversarial success probability is at most $\frac{1}{2} + \frac{1}{2}\sqrt{N\mathsf{Col}(P)}$, where $\mathsf{Col}(P) = \sum_{k\in\mathcal{K}} p_k^2$' i.e., the above scheme satisfies $\frac{1}{2}\sqrt{N\mathsf{Col}(P)}$-unclonable-indistinguishability[11].

The class of unclonable encryption scheme in the above lemma is a generalization of [AS26] in two directions, first, in terms of considering arbitrary hermitian operator family as opposed to Paulis, and second in terms of considering arbitrary distribution on the keyspace as opposed to the uniformly random distribution.

Next we prove Theorem 19 assuming Lemma 9.

Proof of Theorem 19. We will use the interpretation of Lemma 8 and prove $\left(\frac{1}{2}\sqrt{\frac{2^{2n}}{(2^{2n}-1)(2^{2n-1}-1)}}\right)$-unclonable-indistinguishable security of Coset-UE, which by Lemma 8 would imply Theorem 19.

We will show that Coset-UE is an instance of the general class of unclonable indistinguishable encryption schemes in Lemma 9. Let $\mathcal{K} := \{(u,v) \mid u,v \in \mathbb{F}_2^{2n} \setminus \{0\}, \langle u,v\rangle = 0\}$ denote the key space for Coset-UE. For any key $k = (u,v) \in \mathcal{K}$, let $\Theta_k = X^u Z^v \in \mathcal{L}((\mathbb{C}^2)^{\otimes 2n})$. Since for every k, $\Theta_k = X^u Z^v$ is a product of Pauli, it is a unitary and it holds that

$$\Theta_k^\dagger = Z^v X^u = (-1)^{\langle u,v\rangle} X^u Z^v = X^u Z^v = \Theta_k,$$

i.e., Θ_k is Hermitian, and hence, due to unitary condition, $\Theta_k^2 = \Theta_k \Theta_k^\dagger = I$. Moreover, since $(u,v) \neq 0$, $\Theta_k = X^u Z^v$ is a non-identity product of Paulis and therefore $\mathrm{Tr}[\Theta_k] = 0$ and $\Theta_k^2 = I$. Similarly, for any pair of keys, $k = (u,v) \in \mathcal{K}$ and $k' = (u',v')\mathcal{K}$,

$$\mathrm{tr}\left(\Theta_k \Theta_{k'}\right) = \mathrm{tr}\left(X^u Z^v X^{u'} Z^{v'}\right) = \mathrm{tr}\left((-1)^{\langle v,u'\rangle} X^{u+u'} Z^{v+v'}\right) = (-1)^{\langle v,u'\rangle}\,\mathrm{tr}\left(X^{u+u'} Z^{v+v'}\right),$$

which is $\mathrm{tr}(I) = 2^{2n}$ if $u = u', v = v'$, and 0 otherwise. Hence, we conclude that for any pair of keys $k, k' \in \mathcal{K}$, $\mathrm{tr}(\Theta_k \Theta_{k'}) = 2^{2n}\delta_{k,k'}$. Hence $\{\Theta_k\}_{k\in\mathcal{K}}$ satisfies the hypothesis of Lemma 9 with $N = 2^{2n}$.

Next, let P denote the distribution induced by the GenKey algorithm of Coset-UE is just the uniformly random distribution on $\mathcal{K}$. Hence, $\mathsf{Col}(P) = \frac{1}{|\mathcal{K}|}$. It is easy to check that $|K| = (2^{2n} - 1)(2^{2n-1} - 1)$. Hence, we conclude that

$$\mathsf{Col}(P) := \frac{1}{(2^{2n}-1)(2^{2n-1}-1)}. \tag{72}$$

Next we claim that for any message $m \in \{0,1\}$, and key $k \in \mathcal{K}$, the cipher is of the form required in the hypothesis of Lemma 9.

Claim 1. *For any message $m \in \{0,1\}$, and key $k \in \mathcal{K}$, let $\rho_{ct}^{k,m}$ denote the ciphertext produced by Coset-UE.$\mathsf{Enc}(k,m)$, then it holds that*

$$\rho_{ct}^{k,m} = \frac{1}{N}\left(I + (-1)^m \Theta_k\right).$$

Next we complete the proof of the theorem assuming Claim 1 before proving the claim. Let UE be a quantum encryption schme that is the same as Coset-UE, but with a different decryption algorithm, namely the new decryption algorithm Dec′ on input $k, \rho_c t$ simply measure in the eigenbasis

[11] We clarify that for certain choices of the operator family $\{\Theta_k\}_{k\in\mathcal{K}}$ and distribution P, the obtained scheme may not be efficient.

of Θ_k, i.e., the POVM $(\Pi_0, \Pi_1) := \left(\frac{I+\Theta_k}{2}, \frac{I-\Theta_k}{2}\right)$ and outputs the outcome[12]. Clearly, the fact that the operator family $\{\Theta_k\}_{k\in\mathcal{K}}$ satisfies the hypothesis of Lemma 9, combined with Claim 1 imply that UE is an instance of the general class of encryption schemes described in Lemma 9. Hence by Lemma 9, we conclude that UE satisfies $\frac{1}{2}\sqrt{N\mathsf{Col}(P)}$-unclonable indistinguishable security. Since the unclonable-indistinguishable experiment (see Definition 10) is independent of the decryption algorithm of the encryption scheme, and $\mathsf{Coset\text{-}UE}$ and UE only differ in decryption algorithms, $\mathsf{Coset\text{-}UE}$ must also satisfy $\frac{1}{2}\sqrt{N\mathsf{Col}(P)}$-unclonable indistinguishable security. By Equation (72), this implies that $\mathsf{Coset\text{-}UE}$ satisfies

$$\frac{1}{2}\sqrt{N\mathsf{Col}(P)} = \frac{1}{2}\sqrt{\frac{2^{2n}}{(2^{2n}-1)(2^{2n-1}-1)}}\text{-unclonable indistinguishable security.}$$

Next we prove Claim 1 to complete the proof.

Proof of Claim 1. For any $(u,v) \in \mathcal{K}$, let $\text{Compatible}(u,v)$ denote the set of all subspace $A \leq \mathbb{F}_2^{2n}$ with $dim(A) = n$ such that $u \in A$ and $v \in A^\perp$.

For a fixed $A \in \text{Compatible}(u,v)$, let $\text{Coset-rep}(A)$ and $\text{Coset-rep}(A^\perp)$ denote the sets of canonical representatives of the cosets of A and $A^\perp$, respectively, where a canonical representative of a coset is the lexicographically first vector in the coset, and also define

$$\begin{aligned}&\text{Comp-Parity-Coset}((u,v),m,A)\\ &:= \left\{(s,t) \in \text{Coset-rep}(A) \times \text{Coset-rep}(A^\perp) \;\middle|\; \langle v,s\rangle + \langle u,t\rangle = m\right\}.\end{aligned}$$

There are exactly 2^n cosets of A and 2^n cosets of $A^\perp$ and the inner product of every vector in a fixed coset $A+s'$ with v is the same as $\langle v,s\rangle$, and similarly, the inner product of every vector in a fixed coset $A^\perp + t'$ is the same as $\langle u,t\rangle$. Hence, the map $\mathcal{F} : (s+A, t+A^\perp) \longmapsto \langle v,s\rangle + \langle u,t\rangle$ is well-defined on $\left(\mathbb{F}_2^{2n}/A\right) \times \left(\mathbb{F}_2^{2n}/A^\perp\right)$. Since for both u and v, and any bit $b \in \{0,1\}$, there are exactly 2^{2n-1} vectors that have inner product b with u and 2^{2n-1} vectors that have inner product b with v, it holds that $\mathcal{F}$ is a balanced function, i.e., has the same number of preimages of 0 as that for 1. Moreover, for any $m \in \{0,1\}$, the set $\text{Comp-Parity-Coset}((u,v),m,A)$ denotes the preimage set of m under $\mathcal{F}$. Hence

$$|\text{Comp-Parity-Coset}((u,v),m,A)| = \frac{\left|\frac{\mathbb{F}_2^{2n}}{A} \times \frac{\mathbb{F}_2^{2n}}{A^\perp}\right|}{2} = \frac{2^{2n}}{2} = 2^{2n-1}. \tag{73}$$

Since every pair of cosets has the same number of vector representatives, and the parity is constant on each pair of cosets, sampling $s,t \xleftarrow{\$} \mathbb{F}_2^{2n}$ conditioned on $\langle v,s\rangle + \langle u,t\rangle = m$ induces the uniform distribution over $\text{Comp-Parity-Coset}((u,v),m,A)$.

For any fixed key $(u,v) \in \mathcal{K}$, $A \in \text{Compatible}(u,v)$, and message $m \in \{0,1\}$, let $\rho_{ct}^{(u,v),m,A}$ denote the cipher state, conditioned on the sampled subspace being A. Hence, by the above conclusion, we can write

$$\rho_{ct}^{(u,v),m,A} = \frac{1}{|\text{Comp-Parity-Coset}((u,v),m,A)|} \sum_{(s,t)\in\text{Comp-Parity-Coset}((u,v),m,A)} |A_{s,t}\rangle\langle A_{s,t}|.$$

The ciphertext $\rho_{ct}^{(u,v),m} = \mathsf{Coset\text{-}UE.Enc}((u,v),m)$ can then be written as

$$\rho_{ct}^{(u,v),m} = \mathbb{E}_{A\xleftarrow{\$}\text{Compatible}(u,v)} \rho_{ct}^{(u,v),m,A}. \tag{74}$$

[12]Nevertheless, it can be checked that the two decryption algorithms are functionally the same.

Next, we will prove that for any fixed key $(u,v) \in \mathcal{K}$, $A \in \text{Compatible}(u,v)$ and message $m \in \{0,1\}$, it holds that

$$\rho_{ct}^{(u,v),m,A} = \frac{1}{N}\left(I + (-1)^m \Theta_{(u,v)}\right), \tag{75}$$

since that would imply

$$\rho_{ct}^{(u,v),m} = \mathbb{E}_{A \xleftarrow{\$} \text{Compatible}(u,v)} \rho_{ct}^{(u,v),m,A} = \mathbb{E}_{A \xleftarrow{\$} \text{Compatible}(u,v)} \left[\frac{1}{N}\left(I + (-1)^m \Theta_{(u,v)}\right)\right] = \frac{1}{N}\left(I + (-1)^m \Theta_{(u,v)}\right), \tag{76}$$

which proves the claim. Next we prove Equation (75) to complete the proof of the claim. Fix a key $k = (u,v)$, $m \in \{0,1\}$ and $A \in \text{Compatible}(k) = \text{Compatible}(u,v)$ arbitrarily. Since Θ_k is Hermitian, and $\Theta_k^2 = I$, it holds that Θ_k only has eigen values in $\{-1,+1\}$. For any bit $b \in \{0,1\}$, let the eigenspace of $(-1)^b$, i.e., the subspace of eigenvectors with eigenvalue $(-1)^b$ of Θ_k, be denoted as V_b^k. Clearly, the Hilbert space $(\mathbb{C}^2)^{\otimes 2n}$ can be written as the direct sum of V_0^k and V_1^k. Since $\text{Tr}(\Theta_k) = 0$, this implies that $dim(V_0^k) = 2^{2n-1} = dim(V_1^k)$. Moreover, for any bit $b \in \{0,1\}$, the projection on to the eigenspace V_b^k, denoted as $\Pi_{V_b^k}$ can be written as

$$\Pi_{V_b^k} = \frac{1}{2}\left(I + (-1)^b \Theta_k\right). \tag{77}$$

Next, for every $(s,t) \in \text{Comp-Parity-Coset}((u,v),m,A)$, it holds that

$$\begin{aligned}
\Theta_k |A_{s,t}\rangle &= X^u Z^v |A_{s,t}\rangle \\
&= X^u Z^v X^s Z^t |A\rangle && \text{since } X^s Z^t |A\rangle = |A_{s,t}\rangle. \\
&= (-1)^{\langle v,s\rangle} X^u X^s Z^v Z^t |A\rangle && \text{by Pauli anti-commutaion.} \\
&= (-1)^{\langle v,s\rangle} X^{u+s} Z^{v+t} |A\rangle \\
&= (-1)^{\langle v,s\rangle} |A_{u+s,v+t}\rangle \\
&= (-1)^{\langle v,s\rangle}(-1)^{\langle u,t\rangle} |A_{s,t}\rangle \\
&= (-1)^{\langle v,s\rangle + \langle u,t\rangle} |A_{s,t}\rangle = (-1)^m |A_{s,t}\rangle,
\end{aligned}$$

where the last equality follows from the fact that

$$\begin{aligned}
|A_{u+s,v+t}\rangle &= \frac{1}{\sqrt{|A|}} \sum_{a \in A} (-1)^{\langle v+t,a\rangle} |a+s+u\rangle \\
&= \frac{1}{\sqrt{|A|}} \sum_{a \in A} (-1)^{\langle t,a\rangle} |a+s+u\rangle && \text{since } v \in A^\perp. \\
&= \frac{1}{\sqrt{|A|}} \sum_{a' \in A} (-1)^{\langle u,t\rangle}(-1)^{\langle a',t\rangle} |a'+s\rangle && \text{substituting } a' = u+a \text{ and } a = u+a'. \\
&= \frac{(-1)^{\langle u,t\rangle}}{\sqrt{|A|}} \sum_{a' \in A} (-1)^{\langle a',t\rangle} |a'+s\rangle \\
&= (-1)^{\langle u,t\rangle} |A_{s,t}\rangle.
\end{aligned}$$

Therefore the set of vectors $\{|A_{s,t}\rangle\}_{(s,t) \in \text{Comp-Parity-Coset}(k,m,A)} \subset V_m^k$. But clearly, for any two pairs of $(s,t),(s',t') \in \text{Comp-Parity-Coset}(k,m,A)$, such that $(s,t) \neq (s',t')$, $\langle A_{s,t} | A_{s',t'} \rangle = 0$. This follows from the fact that if $s \neq s'$ then they are simply supported on different subsets of computational basis elements, and if $t \neq t'$, then the same holds after applying QFT

on both states, as $QFT\,|A_{s,t}\rangle = \left|A^{\perp}_{t,s}\right\rangle$ and $QFT\left|A_{s',t'}\right\rangle = \left|A^{\perp}_{t',s'}\right\rangle$. Combined with the fact that each element in $\{|A_{s,t}\rangle\}_{(s,t)\in\text{Comp-Parity-Coset}(k,m,A)}$ is a normalized state, we conclude that $\{|A_{s,t}\rangle\}_{(s,t)\in\text{Comp-Parity-Coset}(k,m,A)}$ forms an orthonormal set inside V^k_m. However, $dim(V^k_m) = 2^{2n-1} = |\text{Comp-Parity-Coset}((u,v),m,A)| = |\text{Comp-Parity-Coset}(k,m,A)|$ (see Equation (73)). Therefore it must hold that $\{|A_{s,t}\rangle\}_{(s,t)\in\text{Comp-Parity-Coset}(k,m,A)}$ is an orthonormal basis for V^k_m. Hence, it holds that

$$\Pi_{V^k_m} = \sum_{(s,t)\in\text{Comp-Parity-Coset}(k,m,A)} |A_{s,t}\rangle\langle A_{s,t}|\,. \tag{78}$$

Combined with Equation (77), Equation (78) implies,

$$\begin{aligned}
\rho^{k,m,A}_{ct} &= \frac{1}{|\text{Comp-Parity-Coset}(k,m,A)|} \sum_{(s,t)\in\text{Comp-Parity-Coset}(k,m,A)} |A_{s,t}\rangle\langle A_{s,t}| && (79)\\
&= \frac{1}{2^{2n-1}}\Pi_{V^k_m} && \text{by Equations (73) and (78)} \quad (80)\\
&= \frac{2}{N}\Pi_{V^k_m} && \text{by definition of } N. \quad (81)\\
&= \frac{2}{N}\cdot\frac{1}{2}\left[I + (-1)^m\Theta_k\right] && \text{by Equation (77).} \quad (82)\\
&= \frac{1}{N}\left[I + (-1)^m\Theta_k\right], && (83)
\end{aligned}$$

which completes the proof of Equation (75) and hence the claim. □

□

6.1 Unclonable indistinguishable security of the general class of encryption schemes (Proof of Lemma 9)

We first note that the scheme is perfectly correct. Since Θ_k is a Hermitian operator, with $\Theta_k^2 = I$, all its eigenvalues are in $\{+1,-1\}$. Hence

$$I + (-1)^m\Theta_k \succeq 0.$$

Moreover, since $\text{tr}[\Theta_k] = 0$,

$$\text{tr}\left[\rho^{k,m}_{ct}\right] = \frac{1}{N}\left(\text{tr}[I] + (-1)^m\,\text{tr}[\Theta_k]\right) = 1.$$

Thus $\rho^{k,m}_{ct}$ is a valid density operator. Furthermore, it is supported entirely on the $(-1)^m$-eigenspace of Θ_k, and hence the decryption measurement outputs m with probability one.

Next, we prove unclonable-indistinguishability security. We closely follow the notations and the proof template of [AS26]. In particular, we reprove the Parts 2 and 3 of their main security proof [AS26, Theorem 3.2] for the general case of arbitrary Hermitian operators, and arbitrary distributions on the key space. Next, we recall the following results from [AS26] that will be required for our proof.

Let $\|\|_\infty$ denote the operator norm, i.e., for any operator M, we denote $\|M\|_\infty := \sup_{|\psi|=1}\|M\,|\psi\rangle\|$, where $\|\|$ is the L_2 norm for vectors. Similarly, or any Hermitian operator M, that can be spectral decomposed as $M = \sum_j \lambda_j\,|v_j\rangle\langle v_j|$, we denote $M_+ := \sum_j \max(\lambda_j, 0)\,|v_j\rangle\langle v_j|$.

Theorem 20 (Normalized Choi identity, [AS26, Theorem 2.1]). *Let $\Phi : \mathcal{L}(\mathcal{H}_{\mathsf{E}}) \to \mathcal{L}(\mathcal{H}_{\mathsf{R}_1} \otimes \mathcal{H}_{\mathsf{R}_2})$ be a quantum channel, where $\dim(\mathcal{H}_{\mathsf{E}}) = N$. Let $\mathcal{H}_{\mathsf{A}} \simeq \mathcal{H}_{\mathsf{E}}$, and define*

$$|\Omega\rangle_{\mathsf{AE}} := \frac{1}{\sqrt{N}} \sum_{i=1}^{N} |i\rangle_{\mathsf{A}} |i\rangle_{\mathsf{E}} .$$

Then, for every $M \in \mathcal{L}(\mathcal{H}_{\mathsf{E}})$ and $Q \in \mathcal{L}(\mathcal{H}_{\mathsf{R}_1} \otimes \mathcal{H}_{\mathsf{R}_2})$,

$$\frac{1}{N} \operatorname{tr}[Q\Phi(M)] = \operatorname{tr}\left[(M^T \otimes Q)(I_{\mathsf{A}} \otimes \Phi)(|\Omega\rangle\langle\Omega|)\right],$$

where the transpose is taken in the basis used to define $|\Omega\rangle$.

Lemma 10 (Conditional overlap, [AS26, Lemma 3.3]). *Let $U_{\mathsf{AR}_1} \in \mathcal{L}(\mathcal{H}_{\mathsf{A}} \otimes \mathcal{H}_{\mathsf{R}_1})$ and $V_{\mathsf{AR}_2} \in \mathcal{L}(\mathcal{H}_{\mathsf{A}} \otimes \mathcal{H}_{\mathsf{R}_2})$.*

Viewing them canonically (by tensoring with identity operators) as operators on $\mathcal{H}_{\mathsf{A}} \otimes \mathcal{H}_{\mathsf{R}_1} \otimes \mathcal{H}_{\mathsf{R}_2}$, we have

$$\|UV\|_\infty \le \sqrt{\|\operatorname{tr}_{\mathsf{A}}[U^\dagger U]\|_\infty} \sqrt{\|\operatorname{tr}_{\mathsf{A}}[VV^\dagger]\|_\infty}.$$

Next, we abstract out an operator bound implicit in the proof of [AS26, Proposition 3.8].

Lemma 11 (Abstracted from the Positive-part bound, [AS26, Proposition 3.8]). *Let E_B, E_C, E_{BC} be Hermitian operators on a finite-dimensional Hilbert space, and define the Hermitian operator*

$$G := \frac{1}{2}\left(E_B + E_C + E_{BC} - I\right).$$

Suppose that,

1. *for every $a, c \in \{+1, -1\}$, $\Gamma_{a,c} := \frac{1}{4}\left(I + aE_B + cE_C + acE_{BC}\right) \succeq 0$, and*
2. *there exists $\gamma \ge 0$ such that for every integer $\ell \ge 0$,*

$$\left\| E_B E_{BC}^{\ell} E_C \right\|_\infty \le \gamma.$$

Then, given the spectral decomposition $G = \sum_j \lambda_j |v_j\rangle\langle v_j|$ and $G_+ := \sum_j \max(\lambda_j, 0) |v_j\rangle\langle v_j|$, it holds that

$$\|G_+\|_\infty^2 \le \gamma.$$

We defer the proof of Lemma 11 to Section 6.1.1, which follows from similar arguments as used in the proof of [AS26, Proposition 3.8]. Next we complete the security proof assuming Lemma 11.

Fix an arbitrary adversarial triplet $\mathcal{A}_{\mathsf{UE}} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$.

Let $\Phi : \mathcal{L}((\mathbb{C}^2)^{\otimes \ell}) \to \mathcal{L}(\mathcal{H}_{\mathsf{R}_1} \otimes \mathcal{H}_{\mathsf{R}_2})$ denote the splitting channel applied by $\mathcal{A}$ to produce the registers R_1 and R_2 for $\mathcal{B}$ and $\mathcal{C}$ respectively, before the key is revealed to them.

Next, for every key $k \in \mathcal{K}$, let

$$\{M_{\mathcal{B},0}^k, M_{\mathcal{B},1}^k\} \qquad \text{and} \qquad \{M_{\mathcal{C},0}^k, M_{\mathcal{C},1}^k\}$$

denote the binary POVM corresponding to the measurements applied by $\mathcal{B}$ and $\mathcal{C}$ respectively, after receiving k, and define

$$B_k := M_{\mathcal{B},0}^k - M_{\mathcal{B},1}^k, \qquad C_k := M_{\mathcal{C},0}^k - M_{\mathcal{C},1}^k.$$

It is easy to see that B_k and C_k are Hermitian contractions, and we can rewrite

$$M^k_{\mathcal{B},m} = \frac{I + (-1)^m B_k}{2}, \qquad M^k_{\mathcal{C},m} = \frac{I + (-1)^m C_k}{2}.$$

Let $|\Omega\rangle_{\mathsf{AE}} := \frac{1}{\sqrt{N}} \sum_{i=1}^{N} |i\rangle_{\mathsf{A}} |i\rangle_{\mathsf{E}}$, where E is isomorphic to the ciphertext register A, and let

$$\tau_{\mathsf{AR}_1\mathsf{R}_2} := (I_{\mathsf{A}} \otimes \Phi_{\mathsf{E}}) (|\Omega\rangle\langle\Omega|)$$

be the normalized Choi state of Φ.

For every k, define $\widehat{\Theta}_k := (\Theta_k^T)_{\mathsf{A}}$, and define Define

$$E_B := \sum_{k \in \mathcal{K}} p_k \, \widehat{\Theta}_k \otimes B_k \otimes I, \tag{84}$$

$$E_C := \sum_{k \in \mathcal{K}} p_k \, \widehat{\Theta}_k \otimes I \otimes C_k, \tag{85}$$

$$E_{BC} := \sum_{k \in \mathcal{K}} p_k \, I \otimes B_k \otimes C_k. \tag{86}$$

Also define

$$G := \frac{1}{2} \left(E_B + E_C + E_{BC} - I \right).$$

By Theorem 20,

$$\begin{aligned} &\Pr[b_{\mathcal{B}} = b_{\mathcal{C}} = m] \\ &= \mathbb{E}_{k \leftarrow P,\, m \xleftarrow{\$} \{0,1\}} \operatorname{tr} \left[\left(I + (-1)^m \widehat{\Theta}_k \right) \otimes \frac{I + (-1)^m B_k}{2} \otimes \frac{I + (-1)^m C_k}{2} \, \tau \right]. \end{aligned} \tag{87}$$

Expanding the three factors and averaging over the uniform bit m removes all terms containing an odd power of $(-1)^m$. Hence

$$\begin{aligned} \Pr[b_{\mathcal{B}} = b_{\mathcal{C}} = m] &= \frac{1}{4} \left(1 + \operatorname{tr}[E_B \tau] + \operatorname{tr}[E_C \tau] + \operatorname{tr}[E_{BC} \tau] \right) \\ &= \frac{1}{2} \left(1 + \operatorname{tr}[G \tau] \right) \\ &\leq \frac{1}{2} \left(1 + \|G_+\|_\infty \right), \end{aligned} \tag{88}$$

where the last inequality follows from $G \preceq G_+ \preceq \|G_+\|_\infty I$, and the fact that τ is a density operator.

We now show that the operator bounds used in [AS26] continue to hold when the key is sampled from the possibly nonuniform distribution P.

Let $\gamma := N \sum_{k \in \mathcal{K}} p_k^2 = N \mathsf{Col}(P)$, and define

$$U_{\mathsf{AR}_1} := \sum_{k \in \mathcal{K}} p_k \, \widehat{\Theta}_k \otimes B_k, V_{\mathsf{AR}_2} := \sum_{k \in \mathcal{K}} p_k \, \widehat{\Theta}_k \otimes C_k.$$

Their canonical extensions to the three registers are E_B and E_C, respectively.

Using the Hilbert–Schmidt orthogonality of $\{\Theta_k\}_{k \in \mathcal{K}}$, we obtain

$$\operatorname{tr}_{\mathsf{A}} \left[U^\dagger_{\mathsf{AR}_1} U_{\mathsf{AR}_1} \right] = N \sum_{k \in \mathcal{K}} p_k^2 B_k^2 \preceq \gamma I_{\mathsf{R}_1}, \tag{89}$$

$$\mathrm{tr}_{\mathsf{A}}\left[V_{\mathsf{AR}_2}V_{\mathsf{AR}_2}^{\dagger}\right] = N\sum_{k\in\mathcal{K}} p_k^2 C_k^2 \preceq \gamma I_{\mathsf{R}_2}. \tag{90}$$

Hence, by Lemma 10,

$$\|E_B E_C\|_\infty \le \gamma. \tag{91}$$

We next verify the corresponding extension of [AS26, Proposition 3.6]. For every integer $\ell \ge 0$,

$$\left\|E_B E_{BC}^{\ell} E_C\right\|_\infty \le \gamma. \tag{92}$$

For $\ell = 0$, this follows from Equation (91). For $\ell \ge 1$, expanding E_{BC}^{ℓ} gives

$$E_B E_{BC}^{\ell} E_C = \sum_{k_1,\ldots,k_\ell} p_{k_1}\cdots p_{k_\ell}\left(I\otimes I\otimes C_{k_1}\cdots C_{k_\ell}\right) E_B E_C \left(I\otimes B_{k_1}\cdots B_{k_\ell}\otimes I\right). \tag{93}$$

The two exterior operators are contractions, while the nonnegative coefficients sum to one. Therefore Equation (92) follows from Equation (91).

Finally, for $a, c \in \{+1,-1\}$, define

$$\Gamma_{a,c} := \frac{1}{4}\left(I + aE_B + cE_C + acE_{BC}\right).$$

For every fixed k, the operators $\widehat{\Theta}_k\otimes B_k\otimes I$ and $\widehat{\Theta}_k\otimes I\otimes C_k$ are commuting Hermitian contractions, and their product is $I\otimes B_k\otimes C_k$. Therefore

$$\frac{1}{4}\left(I + a\widehat{\Theta}_k\otimes B_k\otimes I\right)\left(I + c\widehat{\Theta}_k\otimes I\otimes C_k\right) \succeq 0.$$

Averaging with weights p_k shows that

$$\Gamma_{a,c} \succeq 0 \qquad \text{for every } a, c \in \{+1,-1\}. \tag{94}$$

By Equations (92) and (94), the hypotheses of Lemma 11 hold with $\gamma = N\mathsf{Col}(P)$. Therefore

$$\|G_+\|_\infty^2 \le N\mathsf{Col}(P).$$

Combining this with Equation (88) gives $\Pr[b_{\mathcal{B}} = b_{\mathcal{C}} = m] \le \frac{1}{2} + \frac{1}{2}\sqrt{N\mathsf{Col}(P)}$, which completes the proof of the lemma.

6.1.1 Proof of Lemma 11

Proof of Lemma 11. The proof follows the same arguments as in the proof of [AS26, Proposition 3.8].

Let $t := \|G_+\|_\infty$. If $t = 0$, there is nothing to prove. Suppose $t > 0$. Since the Hilbert space is finite dimensional and G is Hermitian, there exists a unit vector $|\psi\rangle$ such that

$$G|\psi\rangle = t|\psi\rangle.$$

First note that the positivity of the four operators $\Gamma_{a,c}$ implies that

$$-I \preceq E_B, E_C, E_{BC} \preceq I.$$

Indeed, suitable pairs of the four operators $\Gamma_{a,c}$ sum to

$$\frac{1}{2}(I \pm E_B), \qquad \frac{1}{2}(I \pm E_C), \qquad \frac{1}{2}(I \pm E_{BC}).$$

Define $\Delta := I - E_{BC}, \qquad R := E_B - E_C$. Then it is easy to see that $\Delta \succeq 0$. Moreover, positivity of $\Gamma_{+1,-1}$ and $\Gamma_{-1,+1}$ gives

$$\Delta + R \succeq 0, \qquad \Delta - R \succeq 0,$$

and hence

$$-\Delta \preceq R \preceq \Delta. \tag{95}$$

Next, the eigenvector equation for G gives

$$(E_B + E_C)\,|\psi\rangle = (2tI + \Delta)\,|\psi\rangle\,. \tag{96}$$

Define $F := 2t(2tI + \Delta)^{-1}$. Since $t > 0$ and $\Delta \succeq 0$, F is well defined and

$$0 \prec F \preceq I, \qquad F(2tI + \Delta) = 2tI.$$

We will first show that

$$RFR \preceq 2t\Delta. \tag{97}$$

For arbitrary vectors $|\xi\rangle, |\eta\rangle$, positivity of $\Delta + R$ and $\Delta - R$ gives

$$\begin{aligned} 0 &\le \frac{1}{2}\langle \xi + \eta|\,(\Delta + R)\,|\xi + \eta\rangle + \frac{1}{2}\langle \xi - \eta|\,(\Delta - R)\,|\xi - \eta\rangle + 2t\|\eta\|^2 \\ &= \langle \xi|\,\Delta\,|\xi\rangle + 2\,\mathrm{Re}\,\langle \xi|\,R\,|\eta\rangle + \langle \eta|\,(\Delta + 2tI)\,|\eta\rangle\,. \end{aligned} \tag{98}$$

Choosing $|\eta\rangle = -(\Delta + 2tI)^{-1}R\,|\xi\rangle$ in Equation (98) yields

$$R(\Delta + 2tI)^{-1}R \preceq \Delta.$$

Multiplying by $2t$ proves Equation (97).

For any vector $|w\rangle$, define

$$Q_F(w) := \langle w|F|w\rangle.$$

Using Equation (96) and $F(2tI + \Delta) = 2tI$, we have

$$Q_F\big((E_B + E_C)\psi\big) = 4t^2 + 2t\langle\psi|\Delta|\psi\rangle. \tag{99}$$

On the other hand, by Equation (97),

$$\begin{aligned} Q_F\big((E_B - E_C)\psi\big) &= \langle\psi|RFR|\psi\rangle \\ &\le 2t\langle\psi|\Delta|\psi\rangle. \end{aligned} \tag{100}$$

Subtracting Equation (100) from Equation (99) and expanding the cross terms gives

$$t^2 \le \mathrm{Re}\langle\psi|E_BFE_C|\psi\rangle \le \|E_BFE_C\|_\infty. \tag{101}$$

It remains to upper bound the last norm. Set $q := (1 + 2t)^{-1} \in (0, 1)$. Since $\Delta = I - E_{BC}$,

$$F = (1 - q)(I - qE_{BC})^{-1}.$$

For an integer $L \geq 1$, define $F_L := (1-q)\sum_{\ell=0}^{L-1} q^\ell E_{BC}^\ell$. By the hypothesis of the lemma,

$$\begin{aligned}\|E_B F_L E_C\|_\infty &\leq (1-q)\sum_{\ell=0}^{L-1} q^\ell \left\|E_B E_{BC}^\ell E_C\right\|_\infty \\ &\leq (1-q)\sum_{\ell=0}^{L-1} q^\ell \gamma \\ &= (1-q^L)\gamma. \end{aligned} \tag{102}$$

Moreover, $F - F_L = q^L E_{BC}^L F$.
As observed above,

$$\|E_B\|_\infty, \|E_C\|_\infty, \|E_{BC}\|_\infty, \|F\|_\infty \leq 1.$$

Therefore

$$\|E_B(F - F_L)E_C\|_\infty = \|E_B(q^L E_{BC}^L F)E_C\|_\infty = q^L \|E_B E_{BC}^L F E_C\|_\infty \leq q^L.$$

Letting $L \to \infty$ in conjunction with Equation (102) gives

$$\|E_B F E_C\|_\infty \leq \gamma.$$

Combining this with Equation (101), we obtain $t^2 \leq \gamma$. Since $t = \|G_+\|_\infty$, this concludes the proof of the lemma. □

7 Computational Monogamy-of-Entanglement Games

In this subsection, we prove the computational versions of the Decisional-Coset-MOE game that will be used in the security proofs of SDE and UPO. The proof follows the computational subspace-hiding proof of [CLLZ21], and its extension to simulatable auxiliary information in [ABH+26, ÇG26]. We deviate from the formalism [ABH+26, ÇG26] who considered independent auxiliary -information generators for $\mathcal{B}$ and $\mathcal{C}$ and instead use a joint auxiliary information generator for $\mathcal{B}$ and $\mathcal{C}$, rather than two independent generators, since the applications in this article allow the two auxiliary strings to be arbitrarily correlated.

For the computational games, let

$$V := \mathbb{F}_2^{4n}.$$

For any subspace $S \leq V$ and $z \in V$, let P_{S+z} denote the canonical circuit that, on input $x \in V$, outputs 1 if and only if $x \in S + z$. All circuits compared using iO below are padded to the same size. We condition on the event that all the generated membership programs are correct on every input; the complement of this event contributes a negligible term.

We first record the following immediate extension of Theorem 19, which will be used at the end of the computational reduction.

Corollary 2 (Post-split auxiliary information)**.** *Consider the Decisional-Coset-MOE experiment, except that after sampling (u, v), the challenger applies an arbitrary, possibly inefficient, classical algorithm*

$$(z_{\mathcal{B}}, z_{\mathcal{C}}) \leftarrow \mathsf{Post}(A, u, v),$$

and sends $(u, v, z_{\mathcal{B}})$ to $\mathcal{B}$ and $(u, v, z_{\mathcal{C}})$ to $\mathcal{C}$. The algorithm Post does not receive s, t or the split registers, and its two outputs may be arbitrarily correlated. Then, for every triplet of adversaries $\mathcal{A}_{moe} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$,

$$\Pr[\mathsf{win}] \leq \frac{1}{2} + \frac{1}{2}\sqrt{\frac{2^{2n}}{(2^{2n}-1)(2^{2n-1}-1)}}.$$

Proof. Fix a nonzero orthogonal pair $k = (u, v)$ and a bit m. By Claim 1, for every compatible subspace A, the state conditioned on k, m, A is

$$\rho_{ct}^{k,m,A} = \frac{1}{2^{2n}}\left(I + (-1)^m \Theta_k\right),$$

which is independent of A. Moreover, for every fixed compatible A and k, the bit

$$m = \langle u, t\rangle + \langle v, s\rangle$$

is uniform. Hence, conditioned on k, the joint distribution of $(A, z_{\mathcal{B}}, z_{\mathcal{C}})$ is independent of m.

The splitting adversary may append a shared classical random string to the two output registers. After receiving k, both recipients can use this common random string to sample the same value of A according to its conditional distribution given k, and then sample the correlated pair $(z_{\mathcal{B}}, z_{\mathcal{C}})$ according to $\mathsf{Post}(A, k)$. This does not require communication and does not use m. Therefore, the post-split auxiliary information can be absorbed into the local strategies of $\mathcal{B}$ and $\mathcal{C}$, and the corollary follows from Theorem 19. □

Next, we define the auxiliary information allowed in the computational game.

Definition 39 (Simulatable pre-split auxiliary information)**.** *Let $\mathsf{Aux} = (\mathit{Aux.Setup}, \mathit{Aux.Gen})$ be a pair of QPT algorithms with the following syntax.*

1. $pp \leftarrow \mathsf{Aux.Setup}(1^\lambda)$: *takes as input the security parameter* 1^λ*, and outputs a classical public parameter* pp*.*

2. $(z_{1,\mathcal{B}}, z_{2,\mathcal{B}}, z_{1,\mathcal{C}}, z_{2,\mathcal{C}}) \leftarrow \mathsf{Aux.Gen}(pp, A, u, v)$: *takes as input a classical public parameter* pp*, a subspace* $A \leq V$*, vectorsu* $\in A \setminus \{0\}$*, and* $v \in A^\perp \setminus \{0\}$*, and outputs pre-split auxiliary information* $(z_{1,\mathcal{B}}, z_{1,\mathcal{C}})$ *and post-split auxiliary information* $(z_{2,\mathcal{B}}, z_{2,\mathcal{C}})$*, for* $\mathcal{B}$ *and* $\mathcal{C}$ *respectively*[13]*.*

There exists a possibly inefficient algorithm $\mathsf{Aux.Sim}$ *such that for every subspace* $A \leq V$*,* $u \in A \setminus \{0\}, v \in A^\perp \setminus \{0\}$*,*

$$\left\{(pp, z_{1,\mathcal{B}}, z_{1,\mathcal{C}}) : \begin{array}{l} pp \leftarrow \mathsf{Aux.Setup}(1^\lambda), \\ (z_{1,\mathcal{B}}, z_{2,\mathcal{B}}, z_{1,\mathcal{C}}, z_{2,\mathcal{C}}) \leftarrow \mathsf{Aux.Gen}(pp, A, u, v) \end{array}\right\} \approx_s \left\{(pp, \widetilde{z}_{1,\mathcal{B}}, \widetilde{z}_{1,\mathcal{C}}) : \begin{array}{l} pp \leftarrow \mathsf{Aux.Setup}(1^\lambda), \\ (\widetilde{z}_{1,\mathcal{B}}, \widetilde{z}_{1,\mathcal{C}}) \leftarrow \mathsf{Aux.Sim}(pp) \end{array}\right\} \tag{103}$$

where the first distribution is obtained by sampling $pp \leftarrow \mathsf{Aux.Setup}(1^\lambda)$ *and then running* $\mathsf{Aux.Gen}(pp, A, u, v)$*, and the second is obtained by sampling* $pp \leftarrow \mathsf{Aux.Setup}(1^\lambda)$ *and then running* $\mathsf{Aux.Sim}(pp)$*.*

Definition 40 (Comp-Decisional-Coset-MOE)**.** *Let* $n := n(\lambda)$ *be a polynomial function in* λ*, and* $V = \mathbb{F}_2^{4n}$*. For a triplet of QPT adversaries* $\mathcal{A}_{moe} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ *and auxiliary algorithms* Aux *as above, consider the following experiment.* $\mathsf{Comp\text{-}Decisional\text{-}Coset\text{-}MOE}_{\mathcal{A}_{moe},\mathsf{Aux}}(1^\lambda)$*:*

1. Ch *samples a uniformly random* $2n$*-dimensional subspace* $A \leq V$ *and independent* $s, t \xleftarrow{\$} V$*, and prepares*

$$|A_{s,t}\rangle = \frac{1}{\sqrt{|A|}} \sum_{a \in A} (-1)^{\langle a,t\rangle} |a + s\rangle .$$

   It also generates $M_0 \leftarrow \mathsf{iO}(P_{A+s})$*, and* $M_1 \leftarrow \mathsf{iO}(P_{A^\perp + t})$*.*

2. Ch *samples* $u \xleftarrow{\$} A \setminus \{0\}, v \xleftarrow{\$} A^\perp \setminus \{0\}$*, and generates* $pp \leftarrow \mathsf{Aux.Setup}(1^\lambda)$*, and samples* $(z_{1,\mathcal{B}}, z_{2,\mathcal{B}}, z_{1,\mathcal{C}}, z_{2,\mathcal{C}}) \leftarrow \mathsf{Aux.Gen}(pp, A, u, v)$*.*

3. Ch *sends* $(|A_{s,t}\rangle, M_0, M_1, pp, z_{1,\mathcal{B}}, z_{1,\mathcal{C}})$ *to* $\mathcal{A}$*.*

4. $\mathcal{A}$ *outputs a joint state* $\sigma_{\mathsf{R}_1,\mathsf{R}_2}$ *and sends registers* R_1 *to* $\mathcal{B}$ *and* R_2 *to* $\mathcal{C}$*.*

5. *After the split,* Ch *sends* $(u, v, z_{2,\mathcal{B}})$ *to* $\mathcal{B}$ *and* $(u, v, z_{2,\mathcal{C}})$ *to* $\mathcal{C}$*.*

6. $\mathcal{B}$ *and* $\mathcal{C}$ *output* $b_\mathcal{B}, b_\mathcal{C}$*.*

7. *The output of the experiment is* 1 *if* $b_\mathcal{B} = b_\mathcal{C} = \langle u, t\rangle + \langle v, s\rangle$*.*

Note that the pair (u, v) is sampled before the split only because the pre-split auxiliary information may depend on it.

For convenience, define

$$\epsilon_{\mathsf{DCMOE},n} := \frac{1}{2}\sqrt{\frac{2^{2n}}{(2^{2n}-1)(2^{2n-1}-1)}}, \qquad \beta_n := \frac{2^n - 1}{2^{2n} - 1}.$$

[13]The four strings $(z_{1,\mathcal{B}}, z_{2,\mathcal{B}}, z_{1,\mathcal{C}}, z_{2,\mathcal{C}})$ can be arbitrarily correlated.

Theorem 21 (Computational Decisional Coset Monogamy with auxiliary information)**.** *Assume post-quantum secure* iO *and post-quantum secure subspace-hiding obfuscation. Then, for every function* $n = n(\lambda) \in \omega(\log(\lambda))$*, and for every QPT triplet* $\mathcal{A}_{moe} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ *and every auxiliary-information generator satisfying Definition 39, there exists a negligible function* $\mathsf{negl}(\lambda)$ *such that*

$$\Pr\left[\textit{Comp-Decisional-Coset-MOE}_{\mathcal{A}_{moe},\mathsf{Aux}}(1^\lambda) = 1\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

Proof. Fix a QPT triplet of adversaries $\mathcal{A}_{\mathsf{moe}} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$. We proceed through a sequence of hybrids, following the computational monogamy proof of [CLLZ21] and the auxiliary-information hybrids of [ABH+26, ÇG26].

Hyb_0**.** This is the experiment $\mathsf{Comp\text{-}Decisional\text{-}Coset\text{-}MOE}_{\mathcal{A}_{\mathsf{moe}},\mathsf{Aux}}(1^\lambda)$.

Hyb_1**: Enlarge the primal subspace.** Replace the first public program $M_0 = \mathsf{iO}(P_{A+s})$ by $\mathsf{iO}(x \mapsto \widehat{U}(x-s))$, where $U \supset A$ is a uniformly random $3n$-dimensional superspace and $\widehat{U} \leftarrow \mathsf{shO}(U)$. Formally, this change is performed through the standard intermediate hybrid in which M_0 is first replaced by $\mathsf{iO}(x \mapsto \widehat{A}(x-s))$ for $\widehat{A} \leftarrow \mathsf{shO}(A)$, followed by the subspace-hiding replacement of $\widehat{A}$ by $\widehat{U}$. The reduction knows (A, s, t, u, v) and can therefore prepare the coset state and generate all four auxiliary strings. Hence Hyb_0 and Hyb_1 are computationally indistinguishable.

Hyb_2**: Rerandomize the primal shift.** Sample $w_s \xleftarrow{\$} U$, set $\sigma := s + w_s$, and replace the source shift s in the first public program by σ. Since $U + s = U + \sigma$, the two source circuits are functionally equivalent, and therefore Hyb_1 and Hyb_2 are computationally indistinguishable by iO-security.

Hyb_3**: Enlarge the dual subspace.** Sample a uniformly random n-dimensional subspace $L \subset A$, so that $L^\perp$ is a uniformly random $3n$-dimensional superspace of $A^\perp$, and replace the second public program $M_1 = \mathsf{iO}(P_{A^\perp + t})$ by $\mathsf{iO}(y \mapsto \widehat{L^\perp}(y-t))$, where $\widehat{L^\perp} \leftarrow \mathsf{shO}(L^\perp)$. As in Hyb_1, this is performed through the standard intermediate hybrid using an obfuscation of $A^\perp$, and subspace-hiding security gives that Hyb_2 and Hyb_3 are computationally indistinguishable.

Hyb_4**: Rerandomize the dual shift.** Sample $w_t \xleftarrow{\$} L^\perp$, set $\tau := t + w_t$, and replace the source shift t in the second public program by τ. Since $L^\perp + t = L^\perp + \tau$, iO-security gives that Hyb_3 and Hyb_4 are computationally indistinguishable.

The first four hybrids are exactly the primal and dual subspace-hiding and affine-rerandomization hybrids used in [CLLZ21]. Their validity is unaffected by the joint auxiliary information because the reduction can sample the auxiliary information as it knows the smaller subspace and all values required to run $\mathsf{Aux.Setup}$ and $\mathsf{Aux.Gen}$. Thus, there exists a negligible function negl such that

$$|\Pr[\mathsf{Hyb}_0 = 1] - \Pr[\mathsf{Hyb}_4 = 1]| \leq \mathsf{negl}(\lambda).$$

Hyb_5**: Simulate the pre-split auxiliary information.** Sample $(\widetilde{z}_{1,\mathcal{B}}, \widetilde{z}_{1,\mathcal{C}}) \leftarrow \mathsf{Aux.Sim}(pp)$ and use this pair in place of $(z_{1,\mathcal{B}}, z_{1,\mathcal{C}})$. Then sample $(z_{2,\mathcal{B}}, z_{2,\mathcal{C}})$ jointly from the conditional distribution given by

$$(z_{1,\mathcal{B}}, z_{2,\mathcal{B}}, z_{1,\mathcal{C}}, z_{2,\mathcal{C}}) \leftarrow \mathsf{Aux.Gen}(pp, A, u, v)),$$

conditioned on the event

$$(z_{1,\mathcal{B}}, z_{1,\mathcal{C}}) = (\widetilde{z}_{1,\mathcal{B}}, \widetilde{z}_{1,\mathcal{C}}).$$

If this event $(z_{1,\mathcal{B}}, z_{1,\mathcal{C}}) = (\widetilde{z}_{1,\mathcal{B}}, \widetilde{z}_{1,\mathcal{C}})$ has probability zero, define the conditional distribution arbitrarily. By Definition 39 and the fact that statistical distance does not increase under a classical channel, there exists a negligible function $\mathsf{negl}'(\lambda)$ such that

$$|\Pr[\mathsf{Hyb}_4 = 1] - \Pr[\mathsf{Hyb}_5 = 1]| \leq \mathsf{negl}'(\lambda).$$

The experiment Hyb_5 may be inefficient, but this is harmless since the transition is statistical and no further computational indistinguishability argument will be used.

In Hyb_5, the entire input given to $\mathcal{A}$ before the split is independent of (u, v) conditioned on (A, s, t) and the two public programs. Hence, the experiment is identically distributed if (u, v) is sampled only after $\mathcal{A}$ outputs the two registers, and $(z_{2,\mathcal{B}}, z_{2,\mathcal{C}})$ is subsequently sampled from the conditional distribution above.

Hyb_6: Reveal L, U, σ, τ. In addition to the information already given before the split, give (L, U, σ, τ) to $\mathcal{A}$. This can only increase the winning probability, and therefore

$$\Pr[\mathsf{Hyb}_5 = 1] \leq \Pr[\mathsf{Hyb}_6 = 1].$$

Analysis of Hyb_6. Condition on fixed values of $(L, U, \sigma, \tau, pp, \widetilde{z}_{1,\mathcal{B}}, \widetilde{z}_{1,\mathcal{C}})$ and on the two public-program strings. Conditioned on these values, A is uniformly distributed among the $2n$-dimensional subspaces satisfying $L \subset A \subset U$. Moreover, $s + \sigma$ is uniform over U, $t + \tau$ is uniform over $L^\perp$, and these two random variables are independent.

Choose an invertible linear transformation T, as a deterministic function of (L, U), such that

$$T(L) = \mathbb{F}_2^n \oplus 0 \oplus 0 \qquad \text{and} \qquad T(U) = \mathbb{F}_2^n \oplus \mathbb{F}_2^{2n} \oplus 0.$$

We use T on computational-basis coordinates and $T^{-\mathsf{T}}$ on phase coordinates. For every intermediate subspace A, there exists a unique n-dimensional subspace $\overline{A} \leq \mathbb{F}_2^{2n}$ such that

$$T(A) = \mathbb{F}_2^n \oplus \overline{A} \oplus 0, \qquad T^{-\mathsf{T}}(A^\perp) = 0 \oplus \overline{A}^\perp \oplus \mathbb{F}_2^n.$$

Since A is uniform subject to $L \subset A \subset U$, the induced subspace $\overline{A}$ is uniformly distributed among the n-dimensional subspaces of $\mathbb{F}_2^{2n}$.

By the same change-of-basis calculation as in [CLLZ21, ABH+26, ÇG26], after applying the known Pauli correction $X^\sigma Z^\tau$ and the unitary corresponding to T, the coset state becomes, up to a global phase,

$$|+\rangle^{\otimes n} \otimes \left|\overline{A}_{\overline{s},\overline{t}}\right\rangle \otimes |0\rangle^{\otimes n},$$

where $\overline{s}, \overline{t} \xleftarrow{\$} \mathbb{F}_2^{2n}$ are independent and uniform. The two outer registers are fixed and can be incorporated into $\mathcal{A}$'s strategy. Write the transformed vectors as

$$T(u) = (u_L, \overline{u}, 0), \qquad T^{-\mathsf{T}}(v) = (0, \overline{v}, v_R),$$

where $\overline{u} \in \overline{A}$ and $\overline{v} \in \overline{A}^\perp$. The target bit satisfies

$$\langle u, t\rangle + \langle v, s\rangle = \langle u, \tau\rangle + \langle v, \sigma\rangle + \langle \overline{u}, \overline{t}\rangle + \langle \overline{v}, \overline{s}\rangle. \tag{104}$$

The first two terms depend only on (σ, τ, u, v) and are therefore known to both recipients after the split. Next, let Bad denote the event that $\overline{u} = 0$ or $\overline{v} = 0$. Since $\overline{u} = 0$ exactly when $u \in L$, and $\overline{v} = 0$ exactly when $v \in U^\perp$, we have

$$\Pr[\overline{u} = 0] = \Pr[\overline{v} = 0] = \frac{2^n - 1}{2^{2n} - 1} = \beta_n.$$

Therefore,

$$\Pr[\mathsf{Bad}] \le 2\beta_n. \tag{105}$$

Conditioned on $\neg\mathsf{Bad}$, the subspace $\overline{A}$ remains uniform and, conditioned on $\overline{A}$, the vectors $\overline{u} \xleftarrow{\$} \overline{A} \setminus \{0\}$ and $\overline{v} \xleftarrow{\$} \overline{A}^{\perp} \setminus \{0\}$ are independent and uniform.

We now reduce Hyb_6, conditioned on the fixed values above and on $\neg\mathsf{Bad}$, to the information-theoretic experiment of Corollary 2. The information-theoretic $\mathcal{A}$ receives $\left|\overline{A}_{\overline{s},\overline{t}}\right\rangle$, appends $|+\rangle^{\otimes n}$ and $|0\rangle^{\otimes n}$, applies the inverse of the known change of basis and Pauli correction, supplies the fixed public programs and the fixed strings $(pp, \widetilde{z}_{1,\mathcal{B}}, \widetilde{z}_{1,\mathcal{C}}, L, U, \sigma, \tau)$, and runs the splitting algorithm from Hyb_6.

After the information-theoretic challenger reveals $(\overline{u}, \overline{v})$, the post-split algorithm samples $u_L, v_R \xleftarrow{\$} \mathbb{F}_2^n$ and defines

$$u := T^{-1}(u_L, \overline{u}, 0), \qquad v := T^{\mathsf{T}}(0, \overline{v}, v_R).$$

These are distributed exactly as the original vectors conditioned on $\neg\mathsf{Bad}$. The algorithm reconstructs

$$A = T^{-1}\left(\mathbb{F}_2^n \oplus \overline{A} \oplus 0\right),$$

samples $(z_{2,\mathcal{B}}, z_{2,\mathcal{C}})$ jointly according to the conditional distribution used in Hyb_5, and computes $\eta := \langle u, \tau\rangle + \langle v, \sigma\rangle$. It gives $(u, v, z_{2,\mathcal{B}}, \eta)$ to $\mathcal{B}$ and $(u, v, z_{2,\mathcal{C}}, \eta)$ to $\mathcal{C}$. This procedure may be inefficient, but it does not receive $(\overline{s}, \overline{t})$ or the split registers and is therefore allowed by Corollary 2.

Each recipient runs the corresponding recipient from Hyb_6. If the recipient outputs b_X, the information-theoretic recipient outputs $b_X + \eta$. By Equation (104), the winning events in the two experiments are identical. Therefore,

$$\Pr[\mathsf{Hyb}_6 = 1 \mid \neg\mathsf{Bad}] \le \frac{1}{2} + \epsilon_{\mathsf{DCMOE},n}.$$

Let $\delta := \Pr[\mathsf{Bad}]$. Upper bounding the winning probability conditioned on Bad by one and using Equation (105), we obtain

$$\begin{aligned}\Pr[\mathsf{Hyb}_6 = 1] &\le \delta + (1-\delta)\left(\frac{1}{2} + \epsilon_{\mathsf{DCMOE},n}\right)\\ &\le \frac{1}{2} + \epsilon_{\mathsf{DCMOE},n} + \frac{\delta}{2}\\ &\le \frac{1}{2} + \epsilon_{\mathsf{DCMOE},n} + \beta_n.\end{aligned}$$

Combining the computational hybrid transitions, the statistical transition from Hyb_4 to Hyb_5, and the inequality between Hyb_5 and Hyb_6, and absorbing all negligible functions into a single negligible function, gives

$$\Pr[\mathsf{Hyb}_0 = 1] \le \frac{1}{2} + \epsilon_{\mathsf{DCMOE},n} + \beta_n + \mathsf{negl}(\lambda).$$

In particular, since $n = \omega(\log\lambda)$, both $\epsilon_{\mathsf{DCMOE},n}$ and β_n are negligible, which completes the proof of the theorem. □

Corollary 3 (Computational Decisional Coset Monogamy without auxiliary information)**.** *Let* $\mathsf{Comp\text{-}Decisional\text{-}Coset\text{-}MOE}_{\mathcal{A}_{moe}}(1^\lambda)$ *denote the experiment of Definition 40 with no auxiliary information, i.e., with*

$$pp = z_{1,\mathcal{B}} = z_{2,\mathcal{B}} = z_{1,\mathcal{C}} = z_{2,\mathcal{C}} = \perp.$$

Assume post-quantum secure iO *and post-quantum secure subspace-hiding obfuscation. Then, for every function* $n = n(\lambda) \in \omega(\log \lambda)$ *and every QPT triplet* $\mathcal{A}_{\mathsf{moe}} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$*, there exists a negligible function* $\mathsf{negl}(\lambda)$ *such that*

$$\Pr\left[\mathit{Comp\text{-}Decisional\text{-}Coset\text{-}MOE}_{\mathcal{A}_{\mathsf{moe}}}(1^\lambda) = 1\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

Proof. Let $\mathsf{Aux.Setup}(1^\lambda)$ output $\bot$, let $\mathsf{Aux.Gen}(pp, A, u, v)$ output $(\bot, \bot, \bot, \bot)$, and let $\mathsf{Aux.Sim}(pp)$ output $(\bot, \bot)$. The statistical simulatability condition in Definition 39 then holds perfectly, and the result follows directly from Theorem 21. □

The subspace-hiding obfuscators required by Corollary 3 can be obtained from post-quantum secure indistinguishability obfuscation and injective keyed one-way functions (see [CLLZ21]) which can be constructed from indistinguishability obfuscation and one-way functions [BPW16].

Part III
Copy-Protection Constructions

8 Single-Decryptor Encryption Construction

In this section, we construct an SDE scheme for the message space $\mathcal{M} = \{0,1\}^{m(\lambda)}$ and prove that it satisfies correlated-challenge decision security. The construction is a minor variant of the coset-state SDE scheme of Kitagawa and Yamakawa ([KY25]): we combine the two ciphertext programs in their construction into a single program with a selector bit.

8.1 Construction

Let $i\mathcal{O}$ be an indistinguishability obfuscation scheme and set $n = n(\lambda) := 2\lambda$. For the ciphertext programs used below, let $\Omega_{i\mathcal{O}} = \{0,1\}^{\ell_{i\mathcal{O}}(\lambda)}$ be the random-tape space of $i\mathcal{O}$. We regard $i\mathcal{O}$ as deterministic once this tape is fixed. All circuits compared in an $i\mathcal{O}$ hybrid are padded to the same size, and $\perp$ is represented by a fixed string distinct from the message outputs.

$\underline{\mathsf{SDE}.\mathsf{GenKey}(1^\lambda)}$

1. Sample a uniformly random n-dimensional subspace $A \leq \mathbb{F}_2^{2n}$ and independent $s, t \xleftarrow{\$} \mathbb{F}_2^{2n}$.
2. Compute $\hat{M}_0 \leftarrow i\mathcal{O}(1^\lambda, M_0)$ and $\hat{M}_1 \leftarrow i\mathcal{O}(1^\lambda, M_1)$, where M_0, M_1 are the following programs.

   $\underline{M_0(v)}$
   Hardcoded: $A + s$
   1. If $v \in A + s$, output 1; otherwise, output 0.

   $\underline{M_1(v)}$
   Hardcoded: $A^\perp + t$
   1. If $v \in A^\perp + t$, output 1; otherwise, output 0.

3. Set $\mathrm{pk} := (\hat{M}_0, \hat{M}_1)$ and $\mathsf{sk} := |A_{s,t}\rangle := \frac{1}{\sqrt{|A|}} \sum_{a \in A} (-1)^{\langle a,t \rangle} |a + s\rangle$. Output $(\mathrm{pk}, \mathsf{sk})$.

$\underline{\mathsf{SDE}.\mathsf{Enc}(\mathrm{pk}, \mathrm{m})}$ The randomness space of encryption is $\mathcal{R} := \mathcal{M} \times \Omega_{i\mathcal{O}}$. On randomness $\tau = (r, \omega) \xleftarrow{\$} \mathcal{R}$, proceed as follows.

1. Parse $\mathrm{pk} = (\hat{M}_0, \hat{M}_1)$.
2. Compute $\hat{P} \leftarrow i\mathcal{O}(1^\lambda, P_{\mathrm{m},r}; \omega)$, where $P_{\mathrm{m},r}$ is the following program.

   $\underline{P_{\mathrm{m},r}(c, v)}$
   Hardcoded: $\hat{M}_0, \hat{M}_1, \mathrm{m}, r$
   1. If $c = 0$ and $\hat{M}_0(v) = 1$, output $\mathrm{m} \oplus r$.

2. If $c = 1$ and $\hat{M}_1(v) = 1$, output r.
3. Otherwise, output $\perp$.

3. Output $\mathrm{ct} := \hat{P}$.

Thus, encryption is fully classical and uses an independent uniform pad and an ordinary fresh $i\mathcal{O}$ tape.

$\underline{\mathsf{SDE}.\mathsf{Dec}(s\mathcal{k}, \mathrm{ct})}$

1. Parse $\hat{P} = \mathrm{ct}$.
2. Run $\hat{P}(0, \cdot)$ on $s\mathcal{k}$ coherently, let y_0 denote the outcome. Then, apply the inverse unitary on $s\mathcal{k}$ to restore the state $s\mathcal{k}$ (as in Lemma 2).
3. Apply $H^{\otimes 2n(\lambda)}$ to $s\mathcal{k}$.
4. Run $\hat{P}(1, \cdot)$ on $s\mathcal{k}$ coherently, let y_1 denote the outcome. Then, again apply the inverse of the last unitary to recover the state $s\mathcal{k}$ (as in Lemma 2).
5. Apply $H^{\otimes 2n(\lambda)}$ to $s\mathcal{k}$.
6. Output $y_0 \oplus y_1$.

Theorem 22 (Correlated-challenge-secure SDE)**.** *Assuming post-quantum secure indistinguishability obfuscation and post-quantum one-way functions, the scheme above is perfectly correct and satisfies correlated-challenge decision security (Definition 22). Consequently, it satisfies strong CPA and* CPA^+ *anti-piracy security.*

8.2 Correctness

Lemma 12. *The above construction is perfectly correct.*

Proof of Lemma 12. Fix an honestly generated key pair and a ciphertext encrypting $\mathrm{m} \in \mathcal{M}$ with pad r. For every $a \in A$, correctness of $i\mathcal{O}$ gives $\hat{P}(0, a + s) = \mathrm{m} \oplus r$. The first evaluation therefore returns $y_0 = \mathrm{m} \oplus r$ with certainty and leaves the coset state unchanged. Moreover, $H^{\otimes 2n} |A_{s,t}\rangle = (-1)^{\langle s,t \rangle} |A^{\perp}_{t,s}\rangle$. Since $\hat{P}(1, a^{\perp} + t) = r$ for every $a^{\perp} \in A^{\perp}$, the second evaluation returns $y_1 = r$ with certainty. Thus decryption outputs $(\mathrm{m} \oplus r) \oplus r = \mathrm{m}$, and the final Hadamard transform restores $|A_{s,t}\rangle$ up to a global phase. □

8.3 Proof of security

Proof of Theorem 22. Fix an admissible sampler Samp and a QPT adversary $\mathcal{A} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$ for the correlated-challenge experiment. Let N be the number of challenge pairs per party output by $\mathcal{A}$; since it is encoded in unary, N is polynomially bounded. For $X \in \{\mathcal{B}, \mathcal{C}\}$ and $i \in [N]$, write $(\mu_{X,i,0}, \mu_{X,i,1})$ for X's i-th pair, identifying the two parties' pairs in identical mode. Set $\delta_{X,i} := \mu_{X,i,0} \oplus \mu_{X,i,1}$ and $\mu_{X,i} := \mu_{X,i,b_X} = \mu_{X,i,0} \oplus b_X \cdot \delta_{X,i}$, where $b \cdot z = z$ if $b = 1$ and $b \cdot z = 0^{m(\lambda)}$ otherwise.

For the ith challenge ciphertext that party $X \in \{\mathcal{B}, \mathcal{C}\}$ receives, let $r_{X,i}$ denote the uniform pad hardcoded in its obfuscated circuit. In separate mode, the pads $(r_{X,i})_{X \in \{\mathcal{B},\mathcal{C}\}, i \in [N]}$ are mutually independent. In identical mode, the challenger samples one pad r_i for each $i \in [N]$, uses it in the single ciphertext for that coordinate, and gives the same ciphertext to both parties; equivalently, $r_{\mathcal{B},i} = r_{\mathcal{C},i} := r_i$.

Hyb_0**: Real experiment.** This is the correlated-challenge experiment of Definition 22, instantiated with the scheme above.

Hyb_1**: Use one uniform challenge bit.** In separate mode, run $\mathsf{Samp}(\mathrm{pk}, \mathsf{IND})$ and denote its output by $(b'_{\mathcal{B}}, b'_{\mathcal{C}})$. Set $\Delta := b'_{\mathcal{B}} \oplus b'_{\mathcal{C}}$. In identical mode, set $\Delta := 0$. Next sample $\widetilde{e} \xleftarrow{\$} \{0,1\}$. Use $b_{\mathcal{B}} := \widetilde{e}$ as Bob's challenge bit and $b_{\mathcal{C}} := \widetilde{e} \oplus \Delta$ as Charlie's challenge bit. All other steps are unchanged.

We show that Hyb_0 and Hyb_1 are identically distributed. Fix pk. For $i, j \in \{0,1\}$, let p_{ij} be the probability that $\mathsf{Samp}(\mathrm{pk}, \mathsf{IND})$ outputs (i, j). The uniform-marginal condition gives $p_{00} = p_{11}$ and $p_{01} = p_{10}$. Consequently, conditioned on any value of Δ in its support, $b'_{\mathcal{B}}$ is uniform and $b'_{\mathcal{C}} = b'_{\mathcal{B}} \oplus \Delta$. Thus $(\widetilde{e}, \widetilde{e} \oplus \Delta)$ in Hyb_1 has the same distribution as $(b'_{\mathcal{B}}, b'_{\mathcal{C}})$ in Hyb_0. In identical mode, the sampler's output is a single uniform bit, which has the same distribution as $\widetilde{e}$. Finally, $\mathcal{A}$ produces its messages and split registers before the sampler is invoked, and the sampler receives only pk and uses fresh coins. The equality therefore holds jointly with the complete output of $\mathcal{A}$.

Hyb_2**: Mask the coset parity to get to the challenge bit.** After sampling (A, s, t), additionally sample $u \xleftarrow{\$} A \setminus \{0\}$, $v \xleftarrow{\$} A^{\perp} \setminus \{0\}$, and $\zeta \xleftarrow{\$} \{0,1\}$. Define $\alpha := \langle v, s \rangle$, $\beta := \langle u, t \rangle$, and $e := \alpha \oplus \beta$, and set $b_{\mathcal{B}} := e \oplus \zeta$ and $b_{\mathcal{C}} := e \oplus \zeta \oplus \Delta$. For every fixed choice of the remaining variables, $e \oplus \zeta$ is uniform. Hence Hyb_1 and Hyb_2 are identically distributed. For later use, define $d_{\mathcal{B}} := \zeta$ and $d_{\mathcal{C}} := \zeta \oplus \Delta$; then $b_X = e \oplus d_X$ for $X \in \{\mathcal{B}, \mathcal{C}\}$. This hybrid changes only the sampling of the challenge bits; the pads $r_{X,i}$ are still sampled as specified above.

Hyb_3**: Write the encryption pads using uniform strings.** We change only how the pads, i.e., $(r_{X,i})_{X \in \{\mathcal{B},\mathcal{C}\},\, i \in [N]}$ from Hyb_2 are sampled. In separate mode, sample mutually independent $q_{X,i} \xleftarrow{\$} \mathcal{M}$ for every $X \in \{\mathcal{B}, \mathcal{C}\}$ and $i \in [N]$. In identical mode, sample independent $q_i \xleftarrow{\$} \mathcal{M}$ for $i \in [N]$ and set $q_{\mathcal{B},i} = q_{\mathcal{C},i} := q_i$. In either mode, define $r_{X,i} := q_{X,i} \oplus \beta \cdot \delta_{X,i}$. Conditioned on all variables sampled before the q's, the string $\beta \cdot \delta_{X,i}$ is fixed, so the map from $q_{X,i}$ to $r_{X,i}$ is a bijection on $\mathcal{M}$. Consequently, in separate mode the pads remain mutually independent and uniform. In identical mode, $\delta_{\mathcal{B},i} = \delta_{\mathcal{C},i}$ and $q_{\mathcal{B},i} = q_{\mathcal{C},i}$, so $r_{\mathcal{B},i} = r_{\mathcal{C},i}$ is a single uniform pad, exactly as in Hyb_2. The transformation is independent across coordinates, and therefore the complete joint distribution of the pads is unchanged. Hence Hyb_2 and Hyb_3 are identically distributed.

This way of sampling the pads prepares the next hybrid. Since $b_X = \alpha \oplus \beta \oplus d_X$, the two occurrences of $\beta \cdot \delta_{X,i}$ cancel in $\mu_{X,i,b_X} \oplus r_{X,i}$, giving $\mu_{X,i,b_X} \oplus r_{X,i} = q_{X,i} \oplus \mu_{X,i,0} \oplus (\alpha \oplus d_X) \cdot \delta_{X,i}$. Thus the first branch of the ciphertext program will use the value α available on $A + s$, while its second branch will use the value β available on $A^{\perp} + t$.

Hyb_4**: Replace the ciphertext programs.** For every X, i, set $\gamma_{X,i} := d_X \cdot \delta_{X,i}$ and replace $P_{\mu_{X,i}, r_{X,i}}$ by the following program $Q_{X,i}$.

$\underline{Q_{X,i}(c, w)}$

Hardcoded: $\hat{M}_0, \hat{M}_1, u, v, q_{X,i}, \gamma_{X,i}, \mu_{X,i,0}, \delta_{X,i}$

1. If $c = 0$ and $\hat{M}_0(w) = 1$, output $q_{X,i} \oplus \mu_{X,i,0} \oplus \gamma_{X,i} \oplus \langle v, w \rangle \cdot \delta_{X,i}$.
2. If $c = 1$ and $\hat{M}_1(w) = 1$, output $q_{X,i} \oplus \langle u, w \rangle \cdot \delta_{X,i}$.

3. Otherwise, output $\bot$.

The old and new programs are functionally equivalent. If $w \in A + s$, then $\langle v, w\rangle = \alpha$, and the first branch returns $q_{X,i} \oplus \mu_{X,i,0} \oplus (d_X \oplus \alpha) \cdot \delta_{X,i} = \mu_{X,i,b_X} \oplus r_{X,i}$. If $w \in A^\bot + t$, then $\langle u, w\rangle = \beta$, and the second branch returns $q_{X,i} \oplus \beta \cdot \delta_{X,i} = r_{X,i}$. Both programs return $\bot$ otherwise.

In separate mode, the $2N$ obfuscations use independent tapes, and hence we change the $2N$ programs inside the obfuscations one at a time using $i\mathcal{O}$ security. In identical mode, there is one common old program and one common new program per coordinate $i \in [N]$, and we change the two identical copies of the obfuscations of the N programs. Since N is polynomial, $|\Pr[\mathsf{Hyb}_3 = 1] - \Pr[\mathsf{Hyb}_4 = 1]| \le \mathsf{negl}(\lambda)$.

We reduce Hyb_4 to computational decisional coset monogamy (Corollary 3), instantiated there with $n = \lambda$. Its challenger samples a 2λ-dimensional subspace of $\mathbb{F}_2^{4\lambda}$, exactly as in our construction, where $n = 2\lambda$.

The splitter $\mathcal{A}'$ receives the coset state and the two obfuscated membership programs for the cosets. It treats the programs as pk and runs $\mathcal{A}$ on $(\mathrm{pk}, |A_{s,t}\rangle)$. It samples $\Delta, \zeta, d_\mathcal{B}, d_\mathcal{C}$, the q's, and the $i\mathcal{O}$ tapes as in Hyb_4, and appends to party X's register pk, d_X and $(\mu_{X,i,0}, \delta_{X,i}, q_{X,i}, \omega_{X,i})_{i\in[N]}$. In identical mode, the parties receive the same q_i and tape ω_i for each coordinate; otherwise, all tapes are independent. These values can be generated before (u, v) are revealed and without knowing (u, v) or e.

Upon receiving (u, v), party X' computes $\gamma_{X,i} = d_X \cdot \delta_{X,i}$, constructs and obfuscates every $Q_{X,i}$ using the tape it received for every $i \in [N]$ from $\mathcal{A}'$, and runs X on the resulting ciphertexts, and obtains $\widetilde{b}_X$. In particular, in the identical mode, the two parties use identical program descriptions and tapes, so they obtain the same obfuscation, exactly matching a single ciphertext copied to both parties. Finally, X' outputs $z_X := \widetilde{b}_X \oplus d_X$. Since $b_X = e \oplus d_X$, we have $z_X = e$ exactly when $\widetilde{b}_X = b_X$. Thus the reduction wins the coset-monogamy game exactly when $\mathcal{A}$ wins Hyb_4. By Corollary 3, $\Pr[\mathsf{Hyb}_4 = 1] \le \frac{1}{2} + \mathsf{negl}(\lambda)$, and hence $\Pr\left[\mathsf{SDE\text{-}CORR}_{\mathbf{SDE},\mathcal{A},\mathsf{Samp}}(1^\lambda) = 1\right] \le \frac{1}{2} + \mathsf{negl}(\lambda)$. Strong CPA and CPA^+ security follow from Theorems 11 and 13. □

9 Unclonable Puncturable Obfuscation Construction

In this section, we prove the following theorem.

Theorem 23 (Correlated UPO from $i\mathcal{O}$ and LWE)**.** *Assume the existence of post-quantum polynomially secure $i\mathcal{O}$ and the quantum hardness of LWE. Then, for every constant $c > 0$ and every polynomial-size keyed circuit class $\mathcal{C} = \{C_k : \mathcal{X}_\lambda \to \mathcal{Y}_\lambda\}_{k\in\mathcal{K}_\lambda}$ where $\mathcal{X} = \{0,1\}^{\ell_{\mathsf{in}}}, \mathcal{Y} := \{0,1\}^{\ell_{\mathsf{out}}}$ and $\ell_{\mathsf{in}} = \ell_{\mathsf{in}}(\lambda) \geq \lambda^c$ for some constant $c > 0$, then there exists a UPO scheme for $\mathcal{C}$ that satisfies λ^c-entropic correlated UPO security. Moreover, for the special case of ID_{U}-generalized UPO security, the LWE assumption can be replaced with post-quantum one-way functions.*

9.1 Construction

Fix an arbitrary constant $c > 0$, with respect to which we want to prove Theorem 23. The UPO scheme that we will use to prove Theorem 23 is the UPO construction of [ABH+26, ÇG26]. Next, we recall the construction of [ABH+26] (also considered for copy protection in [ÇG26], independently) as follows.

Let $n := \lceil \lambda^{c/3} \rceil$, let $V := \mathbb{F}_2^{4n}$, and let $\mathcal{Y}_\lambda = \{0,1\}^{\ell_{\mathsf{out}}(\lambda)}$. Let PPRF be a puncturable PRF with domain $\mathcal{X}_\lambda$ and range $\mathcal{Y}_\lambda$.

$\underline{\mathbf{UPO}.\mathsf{Obf}(1^\lambda, C_k)}$

1. Sample a uniformly random $2n$-dimensional subspace $A \leq V$ and independent $s, t \xleftarrow{\$} V$.
2. Sample $\hat{M}_0 \leftarrow i\mathcal{O}(M_0)$ and $\hat{M}_1 \leftarrow i\mathcal{O}(M_1)$, where M_0, M_1 are the following programs.

> $\underline{M_0(a)}$
> **Hardcoded:** $A + s$
>
> 1. If $a \in A + s$, output 1; otherwise, output 0.

> $\underline{M_1(a)}$
> **Hardcoded:** $A^\perp + t$
>
> 1. If $a \in A^\perp + t$, output 1; otherwise, output 0.

3. Sample $K \leftarrow \mathsf{PPRF}.\mathsf{Gen}(1^\lambda)$.
4. Sample $\hat{Q}_0 \leftarrow i\mathcal{O}(Q_0)$ and $\hat{Q}_1 \leftarrow i\mathcal{O}(Q_1)$, where Q_0, Q_1 are the following programs.

> $\underline{Q_0(a, x)}$
> **Hardcoded:** $\hat{M}_0, C_k, K$
>
> 1. If $\hat{M}_0(a) = 0$, output $\perp$.
> 2. Otherwise, output $F_K(x) \oplus C_k(x)$.

$\hat{Q}_1(a, x)$

Hardcoded: $\hat{M}_1, K$

1. If $\hat{M}_1(a) = 0$, output $\perp$.
2. Otherwise, output $F_K(x)$.

5. Output $\rho_k := \left(|A_{s,t}\rangle, \hat{Q}_0, \hat{Q}_1\right)$.

$\mathsf{UPO}.\mathsf{Eval}(\rho_k, x)$

1. Parse $\rho_k = \left(|A_{s,t}\rangle, \hat{Q}_0, \hat{Q}_1\right)$, and let R denote the register containing $|A_{s,t}\rangle$.
2. Using a reversible implementation of $\hat{Q}_0$, coherently evaluate $\hat{Q}_0(\cdot, x)$ on R into a fresh output register, and measure the output register to obtain y_0.
3. Apply $H^{\otimes 4n}$ to R, obtaining $\left|A^{\perp}_{t,s}\right\rangle$ up to a global phase.
4. Using a reversible implementation of $\hat{Q}_1$, coherently evaluate $\hat{Q}_1(\cdot, x)$ on R into a fresh output register, and measure the output register to obtain y_1.
5. Apply $H^{\otimes 4n}$ to R, restoring $|A_{s,t}\rangle$.
6. Output $(\rho_k, y_0 \oplus y_1)$.

Lemma 13 (Correctness)**.** *The above construction is perfectly correct.*

Proof. Every basis vector in the support of $|A_{s,t}\rangle$ belongs to $A + s$. Hence the first coherent evaluation returns $y_0 = F_K(x) \oplus C_k(x)$ and does not disturb the coset state. After applying $H^{\otimes 4n}$, every basis vector in the support of $\left|A^{\perp}_{t,s}\right\rangle$ belongs to $A^{\perp} + t$, and the second coherent evaluation returns $y_1 = F_K(x)$. The final Hadamard transform restores the input state, so the evaluation algorithm outputs $(\rho_k, C_k(x))$ with certainty. □

9.2 Proof of Security

The proof follows the template used in the proof of [ABH+26, Theorem 6.3]. There are two main differences, first we use a probabilistic argument to rewrite the experiment arbitrary correlated challenge bits, as an experiment with a single challenge bit (see the indistinguishability argument in the indistinguishability proof of Hyb_0 and Hyb_1, which crucially relies on the restriction (Equation (52)) of the UPO sampler definition (Definition 26)), which is suitable for reducing to the computational MOE experiment Comp-Decisional-Coset-MOE, secondly, analogous to the [ABH+26, Lemma 5.3], we prove a new simulation lemma (see Lemma 14) for the correlated UPO security definition that itself allows auxiliary information leakage to the adversaries, (see Definition 27), to show that the pre-split auxiliary information in our final hybrid can be simulated to reduce the hybrid to a valid instance of Comp-Decisional-Coset-MOE.

Let $m := 8n$, and for $u, v \in V$ define $\mathsf{msg}(u, v) := u \,\|\, v \in \{0,1\}^m$. Let $\mathcal{H}$ be an efficiently sampleable universal hash family from $\mathcal{X}_\lambda$ to $\{0,1\}^m$. Let $\mathsf{LF} = (\mathsf{LF.Gen}_{\mathsf{inj}}, \mathsf{LF.Gen}_{\mathsf{loss}})$ be the injective/lossy function family from [ABH+26], with domain $\mathcal{X}_\lambda$, such that every function sampled

in injective mode is injective, the two modes are computationally indistinguishable against QPT algorithms, and every function sampled in lossy mode has image size at most $2^{\ell_{\mathsf{img}}}$, where $\ell_{\mathsf{img}} \leq \lambda^{2c/3}$. We use the puncturable PRF and subspace-hiding obfuscation instantiated in [ABH+26]. For $a \in \{0,1\}$ and $z \in \mathcal{Y}_\lambda$, define $a \cdot z := z$ if $a = 1$, and $a \cdot z := 0^{\ell_{\mathsf{out}}}$ if $a = 0$. Set $\varepsilon := 2^{-\lambda^{2c/3}}$. For all sufficiently large λ, $\lambda^c \geq 2m + 2\ell_{\mathsf{img}} + 2\log(1/\varepsilon) + 3$. For jointly distributed classical random variables X, Z, we use average conditional min-entropy, defined as

$$\widetilde{H}_\infty(X \mid Z) := -\log\left(\sum_z \Pr[Z = z] \max_x \Pr[X = x \mid Z = z]\right). \tag{106}$$

In particular, if Y has support size at most 2^ℓ, then $\widetilde{H}_\infty(X \mid Z, Y) \geq \widetilde{H}_\infty(X \mid Z) - \ell$.

Proof of Theorem 23. We need to prove λ^c-entropic correlated UPO security for **UPO**.

Fix any correlated UPO sampler Samp satisfying $\widetilde{H}_\infty(x_X \mid z_{\mathsf{aux}}) \geq \lambda^c$ for every $X \in \{\mathcal{B}, \mathcal{C}\}$, and a QPT triplet of adversaries $\mathcal{A}_{\mathsf{UPO}} = (\mathcal{A}, \mathcal{B}, \mathcal{C})$. We view $\mathcal{A}$ as a two-stage algorithm with retained state. The pre-split auxiliary string z_{aux} is given to the first stage of $\mathcal{A}$ before it chooses $(k, \mu_\mathcal{B}, \mu_\mathcal{C})$, while $w_\mathcal{B}, w_\mathcal{C}$ are carried unchanged through all the hybrids and are given to $\mathcal{B}, \mathcal{C}$, respectively, after the split. We consider the following set of hybrids.

Hyb_0**: Real experiment.** This is the real experiment $\mathsf{Correlated\text{-}UPO}_{\Pi,\mathcal{A}_{\mathsf{UPO}},\mathsf{Samp},\mathcal{C}}(1^\lambda)$ from Definition 27, instantiated with **UPO** as the construction.

Hyb_1**: Sample the challenge bits using one uniform bit.** Run $(z_{\mathsf{aux}}, x_\mathcal{B}, x_\mathcal{C}, b'_\mathcal{B}, b'_\mathcal{C}, w_\mathcal{B}, w_\mathcal{C}) \leftarrow \mathsf{Samp}(1^\lambda)$, set $\Delta := b'_\mathcal{B} + b'_\mathcal{C}$, discard $b'_\mathcal{B}, b'_\mathcal{C}$, sample $\widetilde{e} \xleftarrow{\$} \{0,1\}$, and set $b_\mathcal{B} := \widetilde{e}$ and $b_\mathcal{C} := \widetilde{e} + \Delta$. The remaining steps are unchanged.

We show that Hyb_0 and Hyb_1 are identically distributed. Fix $(\bar{z}_{\mathsf{aux}}, \bar{x}_\mathcal{B}, \bar{x}_\mathcal{C}, \bar{w}_\mathcal{B}, \bar{w}_\mathcal{C})$ in the support of $(z_{\mathsf{aux}}, x_\mathcal{B}, x_\mathcal{C}, w_\mathcal{B}, w_\mathcal{C})$, and let

$$p_{ij} := \Pr\left[b_\mathcal{B} = i, b_\mathcal{C} = j \,\middle|\, \substack{z_{\mathsf{aux}}=\bar{z}_{\mathsf{aux}},\ x_\mathcal{B}=\bar{x}_\mathcal{B},\ x_\mathcal{C}=\bar{x}_\mathcal{C},\\ w_\mathcal{B}=\bar{w}_\mathcal{B},\ w_\mathcal{C}=\bar{w}_\mathcal{C}}\right].$$

By the condition (Equation (52)) in Definition 26, it is easy to see that $p_{00} = p_{11}$ and $p_{01} = p_{10}$, and hence conditioned on $b_\mathcal{B} \oplus b_\mathcal{C} = b$ for some bit b, the marginal distribution on $b_\mathcal{B}$ as well as $b_\mathcal{C}$ remains uniform. Therefore, we conclude that conditioned on $(z_{\mathsf{aux}}, x_\mathcal{B}, x_\mathcal{C}, w_\mathcal{B}, w_\mathcal{C}, \Delta)$, the bit $b_\mathcal{B}$ is uniform and $b_\mathcal{C} = b_\mathcal{B} + \Delta$. The first-stage output and retained state of $\mathcal{A}$ are obtained from z_{aux} and independent randomness, so adjoining them preserves this identical distribution.

Hyb_2**: Write the common uniform bit as the coset parity plus a uniform mask.** In the execution of Obf, after sampling (A, s, t), sample $u \xleftarrow{\$} A \setminus \{0\}$, $v \xleftarrow{\$} A^\perp \setminus \{0\}$, and $\zeta \xleftarrow{\$} \{0,1\}$, and define $\alpha := \langle v, s\rangle$, $\beta := \langle u, t\rangle$, and $e := \alpha + \beta$. Set $b_\mathcal{B} := e + \zeta$ and $b_\mathcal{C} := e + \zeta + \Delta$. Since $e + \zeta$ is uniform for every fixed value of the remaining variables, Hyb_1 and Hyb_2 are identically distributed. Define $d_\mathcal{B} := \zeta$, $d_\mathcal{C} := \zeta + \Delta$, and $\mathsf{diff}_X(x) := C_k(x) \oplus \mu_X(x)$ for $X \in \{\mathcal{B}, \mathcal{C}\}$. Then $b_X = \alpha + \beta + d_X$.

Hyb_3**: Puncture the PRF and hardcode its values at the challenge points.** Let $\mathcal{S} := \{x_\mathcal{B}, x_\mathcal{C}\}$, and let $K_\mathcal{S}$ be the puncturable PRF key punctured at $\mathcal{S}$. If $x_\mathcal{B} \neq x_\mathcal{C}$, let $y_X := F_K(x_X)$ for every $X \in \{\mathcal{B}, \mathcal{C}\}$. If $x_\mathcal{B} = x_\mathcal{C}$, let $y_\mathcal{B} := F_K(x_\mathcal{B})$ and set $y_\mathcal{C} := y_\mathcal{B}$. Replace Q_0, Q_1 by the following programs P_0, P_1.

$\underline{P_0(a,x)}$

Hardcoded: $\hat{M}_0, C_k, \mu_{\mathcal{B}}, \mu_{\mathcal{C}}, K_{\mathcal{S}}, x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}, y_{\mathcal{B}}, y_{\mathcal{C}}$

1. If $\hat{M}_0(a) = 0$, output $\perp$.
2. If $x = x_{\mathcal{B}}$, output $y_{\mathcal{B}} \oplus C_k(x) \oplus b_{\mathcal{B}} \cdot \mathsf{diff}_{\mathcal{B}}(x)$.
3. If $x = x_{\mathcal{C}}$ and $x_{\mathcal{C}} \neq x_{\mathcal{B}}$, output $y_{\mathcal{C}} \oplus C_k(x) \oplus b_{\mathcal{C}} \cdot \mathsf{diff}_{\mathcal{C}}(x)$.
4. Output $F_{K_{\mathcal{S}}}(x) \oplus C_k(x)$.

$\underline{P_1(a,x)}$

Hardcoded: $\hat{M}_1, K_{\mathcal{S}}, x_{\mathcal{B}}, x_{\mathcal{C}}, y_{\mathcal{B}}, y_{\mathcal{C}}$

1. If $\hat{M}_1(a) = 0$, output $\perp$.
2. If $x = x_{\mathcal{B}}$, output $y_{\mathcal{B}}$.
3. If $x = x_{\mathcal{C}}$ and $x_{\mathcal{C}} \neq x_{\mathcal{B}}$, output $y_{\mathcal{C}}$.
4. Output $F_{K_{\mathcal{S}}}(x)$.

The relevant circuits in Hyb_2 and Hyb_3 have the same functionality. At $x = x_{\mathcal{B}}$, the programs in Hyb_3 output $y_{\mathcal{B}} \oplus C_k(x_{\mathcal{B}}) \oplus b_{\mathcal{B}} \cdot \mathsf{diff}_{\mathcal{B}}(x_{\mathcal{B}})$ and $y_{\mathcal{B}}$, respectively, which are exactly the corresponding outputs of the programs in Hyb_2. This also covers the case $x_{\mathcal{B}} = x_{\mathcal{C}}$, since the Bob line is evaluated first. If $x_{\mathcal{B}} \neq x_{\mathcal{C}}$, the same argument applies at $x = x_{\mathcal{C}}$. Finally, for $x \notin \mathcal{S}$, punctured correctness gives $F_{K_{\mathcal{S}}}(x) = F_K(x)$. Therefore, by $i\mathcal{O}$-security, we conclude that Hyb_2 and Hyb_3 are computationally indistinguishable.

Hyb_4**: Replace the hardcoded PRF values by uniform strings.** If $x_{\mathcal{B}} \neq x_{\mathcal{C}}$, replace $(F_K(x_{\mathcal{B}}), F_K(x_{\mathcal{C}}))$ by independent $y_{\mathcal{B}}, y_{\mathcal{C}} \xleftarrow{\$} \mathcal{Y}_\lambda$. If $x_{\mathcal{B}} = x_{\mathcal{C}}$, sample $y_{\mathcal{B}} \xleftarrow{\$} \mathcal{Y}_\lambda$ and set $y_{\mathcal{C}} := y_{\mathcal{B}}$. Puncturable-PRF security implies that Hyb_3 and Hyb_4 are computationally indistinguishable.

Hyb_5**: Rewrite the uniform values.** Sample independent $q_{\mathcal{B}}, q_{\mathcal{C}} \xleftarrow{\$} \mathcal{Y}_\lambda$. If $x_{\mathcal{B}} \neq x_{\mathcal{C}}$, then for every $X \in \{\mathcal{B}, \mathcal{C}\}$ set

$$y_X := q_X \oplus \beta \cdot \mathsf{diff}_X(x_X). \tag{107}$$

If $x_{\mathcal{B}} = x_{\mathcal{C}}$, set

$$y_{\mathcal{B}} = y_{\mathcal{C}} := q_{\mathcal{B}} \oplus \beta \cdot \mathsf{diff}_{\mathcal{B}}(x_{\mathcal{B}}), \tag{108}$$

and let $q_{\mathcal{C}}$ be an unused independent uniform string. Since the values y_X in Hyb_4 are uniform, this is an exact change of variables. Replace P_0, P_1 by the following programs.

$\underline{P'_0(a,x)}$

Hardcoded: $\hat{M}_0, C_k, \mu_{\mathcal{B}}, \mu_{\mathcal{C}}, K_{\mathcal{S}}, x_{\mathcal{B}}, x_{\mathcal{C}}, q_{\mathcal{B}}, q_{\mathcal{C}}, \alpha, d_{\mathcal{B}}, d_{\mathcal{C}}$

1. If $\hat{M}_0(a) = 0$, output $\perp$.
2. If $x = x_{\mathcal{B}}$, output $q_{\mathcal{B}} \oplus C_k(x) \oplus (\alpha + d_{\mathcal{B}}) \cdot \mathsf{diff}_{\mathcal{B}}(x)$.

3. If $x = x_{\mathcal{C}}$ and $x_{\mathcal{C}} \neq x_{\mathcal{B}}$, output $q_{\mathcal{C}} \oplus C_k(x) \oplus (\alpha + d_{\mathcal{C}}) \cdot \mathsf{diff}_{\mathcal{C}}(x)$.

4. Output $F_{K_{\mathcal{S}}}(x) \oplus C_k(x)$.

$\underline{P'_1(a,x)}$

Hardcoded: $\hat{M}_1, C_k, \mu_{\mathcal{B}}, \mu_{\mathcal{C}}, K_{\mathcal{S}}, x_{\mathcal{B}}, x_{\mathcal{C}}, q_{\mathcal{B}}, q_{\mathcal{C}}, \beta$

1. If $\hat{M}_1(a) = 0$, output $\perp$.

2. If $x = x_{\mathcal{B}}$, output $q_{\mathcal{B}} \oplus \beta \cdot \mathsf{diff}_{\mathcal{B}}(x)$.

3. If $x = x_{\mathcal{C}}$ and $x_{\mathcal{C}} \neq x_{\mathcal{B}}$, output $q_{\mathcal{C}} \oplus \beta \cdot \mathsf{diff}_{\mathcal{C}}(x)$.

4. Output $F_{K_{\mathcal{S}}}(x)$.

The relevant circuits in Hyb_4 and Hyb_5 have the same functionality. Suppose first that $x_{\mathcal{B}} \neq x_{\mathcal{C}}$ and fix $X \in \{\mathcal{B}, \mathcal{C}\}$. By the definition of y_X,

$$q_X \oplus \beta \cdot \mathsf{diff}_X(x_X) = y_X. \tag{109}$$

Moreover, since $b_X = \alpha + \beta + d_X$,

$$\begin{aligned} & q_X \oplus C_k(x_X) \oplus (\alpha + d_X) \cdot \mathsf{diff}_X(x_X) \\ &\quad = y_X \oplus C_k(x_X) \oplus (\alpha + \beta + d_X) \cdot \mathsf{diff}_X(x_X) \\ &\quad = y_X \oplus C_k(x_X) \oplus b_X \cdot \mathsf{diff}_X(x_X). \end{aligned}$$

Thus, the individual outputs of P'_0, P'_1 at x_X are exactly the corresponding outputs of P_0, P_1. If $x_{\mathcal{B}} = x_{\mathcal{C}}$, the Bob line is evaluated first, and the same calculation applies with $X = \mathcal{B}$. At every other point, the programs are unchanged. Therefore, by $i\mathcal{O}$-security we conclude that Hyb_4 and Hyb_5 are computationally indistinguishable.

Hyb_6**: Replace the explicit point comparisons by injective-function comparisons.** For every $X \in \{\mathcal{B}, \mathcal{C}\}$, sample $h_X \xleftarrow{\$} \mathcal{H}$ and $L_X \leftarrow \mathsf{LF.Gen}_{\mathsf{inj}}(1^\lambda)$ independently, and set $\eta_X := L_X(x_X)$ and $\mathsf{str}_X := \mathsf{msg}(u, v) \oplus h_X(x_X)$. Replace P'_0, P'_1 by the following programs R_0, R_1.

$\underline{R_0(a,x)}$

Hardcoded: $\hat{M}_0, C_k, \mu_{\mathcal{B}}, \mu_{\mathcal{C}}, K_{\mathcal{S}}, q_{\mathcal{B}}, q_{\mathcal{C}}, d_{\mathcal{B}}, d_{\mathcal{C}}, L_{\mathcal{B}}, L_{\mathcal{C}}, \eta_{\mathcal{B}}, \eta_{\mathcal{C}}, h_{\mathcal{B}}, h_{\mathcal{C}}, \mathsf{str}_{\mathcal{B}}, \mathsf{str}_{\mathcal{C}}$

1. If $\hat{M}_0(a) = 0$, output $\perp$.

2. If $L_{\mathcal{B}}(x) = \eta_{\mathcal{B}}$, parse $(u'_{\mathcal{B}}, v'_{\mathcal{B}}) := \mathsf{str}_{\mathcal{B}} \oplus h_{\mathcal{B}}(x)$ and output

$$q_{\mathcal{B}} \oplus C_k(x) \oplus \big(\langle v'_{\mathcal{B}}, a\rangle + d_{\mathcal{B}}\big) \cdot \mathsf{diff}_{\mathcal{B}}(x).$$

3. If $L_{\mathcal{C}}(x) = \eta_{\mathcal{C}}$, parse $(u'_{\mathcal{C}}, v'_{\mathcal{C}}) := \mathsf{str}_{\mathcal{C}} \oplus h_{\mathcal{C}}(x)$ and output

$$q_{\mathcal{C}} \oplus C_k(x) \oplus \big(\langle v'_{\mathcal{C}}, a\rangle + d_{\mathcal{C}}\big) \cdot \mathsf{diff}_{\mathcal{C}}(x).$$

4. Output $F_{K_{\mathcal{S}}}(x) \oplus C_k(x)$.

$\underline{R_1(a,x)}$

Hardcoded: $\hat{M}_1, C_k, \mu_{\mathcal{B}}, \mu_{\mathcal{C}}, K_{\mathcal{S}}, q_{\mathcal{B}}, q_{\mathcal{C}}, L_{\mathcal{B}}, L_{\mathcal{C}}, \eta_{\mathcal{B}}, \eta_{\mathcal{C}}, h_{\mathcal{B}}, h_{\mathcal{C}}, \mathsf{str}_{\mathcal{B}}, \mathsf{str}_{\mathcal{C}}$

1. If $\hat{M}_1(a) = 0$, output $\perp$.
2. If $L_{\mathcal{B}}(x) = \eta_{\mathcal{B}}$, parse $(u'_{\mathcal{B}}, v'_{\mathcal{B}}) := \mathsf{str}_{\mathcal{B}} \oplus h_{\mathcal{B}}(x)$ and output
$$q_{\mathcal{B}} \oplus \langle u'_{\mathcal{B}}, a\rangle \cdot \mathsf{diff}_{\mathcal{B}}(x).$$
3. If $L_{\mathcal{C}}(x) = \eta_{\mathcal{C}}$, parse $(u'_{\mathcal{C}}, v'_{\mathcal{C}}) := \mathsf{str}_{\mathcal{C}} \oplus h_{\mathcal{C}}(x)$ and output
$$q_{\mathcal{C}} \oplus \langle u'_{\mathcal{C}}, a\rangle \cdot \mathsf{diff}_{\mathcal{C}}(x).$$
4. Output $F_{K_{\mathcal{S}}}(x)$.

The relevant circuits in Hyb_5 and Hyb_6 have the same functionality. Suppose first that $x_{\mathcal{B}} \neq x_{\mathcal{C}}$. At $x = x_X$, the equality $L_X(x_X) = \eta_X$ holds, while the equality corresponding to the other recipient does not hold by injectivity. Moreover, $\mathsf{str}_X \oplus h_X(x_X) = u\|v$. For every $a \in A + s$, it holds that $\langle v, a\rangle = \langle v, s\rangle = \alpha$, and for every $a \in A^{\perp} + t$, it holds that $\langle u, a\rangle = \langle u, t\rangle = \beta$. Therefore, the individual outputs of R_0 and R_1 at x_X are exactly the corresponding outputs of P'_0 and P'_1.

If $x_{\mathcal{B}} = x_{\mathcal{C}}$, both equalities hold at the common point, but the Bob equality is evaluated first in both programs. Hence, the programs use the Bob values independently of Δ, exactly as prescribed by Equation (53). At every $x \notin \mathcal{S}$, neither equality holds, and the final lines of the programs are evaluated. Thus, the two pairs of programs are functionally equivalent. Therefore, by $i\mathcal{O}$-security, we conclude that Hyb_5 and Hyb_6 are computationally indistinguishable.

Hyb_7: Replace the punctured PRF key with the unpunctured key. Replace $K_{\mathcal{S}}$ by K in the final line of R_0, R_1. The relevant circuits have the same functionality, since one of the preceding equality tests holds at every point in $\mathcal{S}$, and $F_{K_{\mathcal{S}}}(x) = F_K(x)$ for every $x \notin \mathcal{S}$. Therefore, by $i\mathcal{O}$-security, we conclude that Hyb_6 and Hyb_7 are computationally indistinguishable.

Hyb_8: Switch the two functions to lossy mode. Replace $L_{\mathcal{B}}$ and $L_{\mathcal{C}}$, one at a time, by functions sampled using $\mathsf{LF.Gen}_{\mathsf{loss}}(1^\lambda)$, and recompute $\eta_X := L_X(x_X)$ after each replacement. Injective/lossy mode indistinguishability implies that Hyb_7 and Hyb_8 are computationally indistinguishable.

Define the auxiliary-information generator (see Definition 39) as follows. The algorithm $\mathsf{Aux.Setup}(1^\lambda)$ samples $h_{\mathcal{B}}, h_{\mathcal{C}} \xleftarrow{\$} \mathcal{H}$ and independent $L_{\mathcal{B}}, L_{\mathcal{C}} \leftarrow \mathsf{LF.Gen}_{\mathsf{loss}}(1^\lambda)$, and outputs $pp := (L_{\mathcal{B}}, L_{\mathcal{C}}, h_{\mathcal{B}}, h_{\mathcal{C}})$. On input (pp, A, u, v), the algorithm $\mathsf{Aux.Gen}$ runs $(z_{\mathsf{aux}}, x_{\mathcal{B}}, x_{\mathcal{C}}, b'_{\mathcal{B}}, b'_{\mathcal{C}}, w_{\mathcal{B}}, w_{\mathcal{C}}) \leftarrow \mathsf{Samp}(1^\lambda)$, sets $\Delta := b'_{\mathcal{B}} + b'_{\mathcal{C}}$, computes $\eta_X := L_X(x_X)$ and $\mathsf{str}_X := \mathsf{msg}(u, v) \oplus h_X(x_X)$ for $X \in \{\mathcal{B}, \mathcal{C}\}$, and outputs

$$z_{1,\mathcal{B}} := (z_{\mathsf{aux}}, \Delta, \eta_{\mathcal{B}}, \eta_{\mathcal{C}}, \mathsf{str}_{\mathcal{B}}, \mathsf{str}_{\mathcal{C}}), \qquad z_{2,\mathcal{B}} := (x_{\mathcal{B}}, w_{\mathcal{B}}),$$
$$z_{1,\mathcal{C}} := \perp, \qquad z_{2,\mathcal{C}} := (x_{\mathcal{C}}, w_{\mathcal{C}}).$$

The simulator $\mathsf{Aux.Sim}(pp)$ runs $(z_{\mathsf{aux}}, x_{\mathcal{B}}, x_{\mathcal{C}}, b'_{\mathcal{B}}, b'_{\mathcal{C}}, w_{\mathcal{B}}, w_{\mathcal{C}}) \leftarrow \mathsf{Samp}(1^\lambda)$, computes $\Delta, \eta_{\mathcal{B}}, \eta_{\mathcal{C}}$, discards $w_{\mathcal{B}}, w_{\mathcal{C}}$, samples independent $U_{\mathcal{B}}, U_{\mathcal{C}} \xleftarrow{\$} \{0,1\}^m$, and outputs $\widetilde{z}_{1,\mathcal{B}} := (z_{\mathsf{aux}}, \Delta, \eta_{\mathcal{B}}, \eta_{\mathcal{C}}, U_{\mathcal{B}}, U_{\mathcal{C}})$

and $\widetilde{z}_{1,\mathcal{C}} := \perp$. Next we show the statistical simulatability condition for the above auxiliary-information generator, as requried in Definition 39.

Lemma 14 (Statistical simulation of the pre-split auxiliary information)**.** *For every fixed subspace $A \le V$, $u \in A \setminus \{0\}$, and $v \in A^{\perp} \setminus \{0\}$, the auxiliary-information algorithms defined above satisfy the statistical simulatability condition in Equation (103) with statistical distance at most ε.*

Proof of Lemma 14. Let $(Z_{\mathsf{aux}}, X_{\mathcal{B}}, X_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}, W_{\mathcal{B}}, W_{\mathcal{C}}) \leftarrow \mathsf{Samp}(1^{\lambda})$ and set $\Delta := b_{\mathcal{B}} + b_{\mathcal{C}}$. The variables $W_{\mathcal{B}}, W_{\mathcal{C}}$ are carried as post-split auxiliary information and do not occur in the pre-split strings. Condition on the descriptions of $L_{\mathcal{B}}, L_{\mathcal{C}}$, set $\eta_X := L_X(X_X)$ for every $X \in \{\mathcal{B}, \mathcal{C}\}$, and let $Z := (\Delta, \eta_{\mathcal{B}}, \eta_{\mathcal{C}})$. Since the descriptions of $L_{\mathcal{B}}, L_{\mathcal{C}}$ are independent of $(Z_{\mathsf{aux}}, X_{\mathcal{B}}, X_{\mathcal{C}}, \Delta)$ and $|\operatorname{Supp}(Z)| \le 2^{1+2\ell_{\mathsf{img}}}$, the support-size leakage bound for average conditional min-entropy gives $\widetilde{H}_{\infty}(X_X \mid Z_{\mathsf{aux}}, Z) \ge \lambda^c - 1 - 2\ell_{\mathsf{img}}$ for every $X \in \{\mathcal{B}, \mathcal{C}\}$.

For every fixed $(Z_{\mathsf{aux}}, Z) = (z_{\mathsf{aux}}, z)$, let $\mathrm{cp}_{\mathcal{B}}(z_{\mathsf{aux}}, z)$, $\mathrm{cp}_{\mathcal{C}}(z_{\mathsf{aux}}, z)$, and $\mathrm{cp}_{\mathcal{B},\mathcal{C}}(z_{\mathsf{aux}}, z)$ denote the collision probabilities of $X_{\mathcal{B}}$, $X_{\mathcal{C}}$, and $(X_{\mathcal{B}}, X_{\mathcal{C}})$, respectively, conditioned on $(Z_{\mathsf{aux}}, Z) = (z_{\mathsf{aux}}, z)$. For every $X \in \{\mathcal{B}, \mathcal{C}\}$,

$$\begin{aligned}\mathbb{E}_{Z_{\mathsf{aux}},Z}[\mathrm{cp}_X(Z_{\mathsf{aux}}, Z)] &\le \sum_{z_{\mathsf{aux}},z} \Pr[Z_{\mathsf{aux}} = z_{\mathsf{aux}}, Z = z] \max_x \Pr[X_X = x \mid Z_{\mathsf{aux}} = z_{\mathsf{aux}}, Z = z] \\ &= 2^{-\widetilde{H}_{\infty}(X_X | Z_{\mathsf{aux}}, Z)} \le 2^{-(\lambda^c - 1 - 2\ell_{\mathsf{img}})}.\end{aligned}$$

Moreover, $\mathrm{cp}_{\mathcal{B},\mathcal{C}}(z_{\mathsf{aux}}, z) \le \min\{\mathrm{cp}_{\mathcal{B}}(z_{\mathsf{aux}}, z), \mathrm{cp}_{\mathcal{C}}(z_{\mathsf{aux}}, z)\}$ for every (z_{aux}, z).

Independence and universality of $h_{\mathcal{B}}, h_{\mathcal{C}}$ give, for every fixed $(Z_{\mathsf{aux}}, Z) = (z_{\mathsf{aux}}, z)$,

$$\begin{aligned}&\mathbb{E}_{h_{\mathcal{B}},h_{\mathcal{C}}}\left[\mathrm{cp}\big(h_{\mathcal{B}}(X_{\mathcal{B}}), h_{\mathcal{C}}(X_{\mathcal{C}})\big) \mid Z_{\mathsf{aux}} = z_{\mathsf{aux}}, Z = z\right] \\ &\quad \le 2^{-2m} + 2^{-m}\big(\mathrm{cp}_{\mathcal{B}}(z_{\mathsf{aux}}, z) + \mathrm{cp}_{\mathcal{C}}(z_{\mathsf{aux}}, z)\big) + \mathrm{cp}_{\mathcal{B},\mathcal{C}}(z_{\mathsf{aux}}, z).\end{aligned}$$

Averaging over (Z_{aux}, Z) therefore gives

$$\begin{aligned}&\mathbb{E}_{Z_{\mathsf{aux}},Z,h_{\mathcal{B}},h_{\mathcal{C}}}\left[\mathrm{cp}\big(h_{\mathcal{B}}(X_{\mathcal{B}}), h_{\mathcal{C}}(X_{\mathcal{C}})\big)\right] \\ &\quad \le 2^{-2m} + \left(1 + 2^{1-m}\right) 2^{-(\lambda^c - 1 - 2\ell_{\mathsf{img}})}.\end{aligned}$$

The collision-probability bound on statistical distance, followed by Jensen's inequality, gives statistical distance at most

$$\frac{1}{2}\sqrt{2^{2m-(\lambda^c-1-2\ell_{\mathsf{img}})}\left(1 + 2^{1-m}\right)} \le \varepsilon, \tag{110}$$

where the last inequality follows from the choice of parameters. Finally, XORing both hash values with the fixed string $\mathsf{msg}(u, v)$ does not change statistical distance. Hence the real and simulated pre-split auxiliary-information distributions satisfy the statistical simulatability condition in Definition 39. □

We construct a QPT triplet of adversaries $\widehat{\mathcal{A}_{\mathsf{moe}}} = (\widehat{\mathcal{A}}, \widehat{\mathcal{B}}, \widehat{\mathcal{C}})$ for Theorem 21. The algorithm $\widehat{\mathcal{A}}$ receives $|A_{s,t}\rangle$, $\hat{M}_0, \hat{M}_1$, pp, and the pre-split auxiliary information, from which it parses $z_{\mathsf{aux}}, \Delta, \eta_{\mathcal{B}}, \eta_{\mathcal{C}}, \mathsf{str}_{\mathcal{B}}, \mathsf{str}_{\mathcal{C}}$. It runs the first stage of $\mathcal{A}$ on $(1^{\lambda}, z_{\mathsf{aux}})$ to obtain $(k, \mu_{\mathcal{B}}, \mu_{\mathcal{C}})$ and retained state, samples $\zeta \xleftarrow{\$} \{0,1\}$, $q_{\mathcal{B}}, q_{\mathcal{C}} \xleftarrow{\$} \mathcal{Y}_{\lambda}$, and $K \leftarrow \mathsf{PPRF.Gen}(1^{\lambda})$, sets $d_{\mathcal{B}} := \zeta$ and $d_{\mathcal{C}} := \zeta + \Delta$, constructs the two lossy-mode programs from Hyb_8, obfuscates them using $i\mathcal{O}$, and gives $\rho_k := (|A_{s,t}\rangle, \hat{R}_0, \hat{R}_1)$ to the continuation of $\mathcal{A}$ together with its retained state. When $\mathcal{A}$ outputs its two registers, $\widehat{\mathcal{A}}$ appends a classical copy of ζ to the register sent to $\widehat{\mathcal{B}}$ and classical copies of (ζ, Δ) to the register sent to $\widehat{\mathcal{C}}$.

After the split, $\widehat{\mathcal{B}}$ receives $(u, v, x_{\mathcal{B}}, w_{\mathcal{B}})$, runs $\mathcal{B}$ on $(x_{\mathcal{B}}, w_{\mathcal{B}})$, and outputs $\widehat{b}_{\mathcal{B}} + \zeta$. The algorithm $\widehat{\mathcal{C}}$ receives $(u, v, x_{\mathcal{C}}, w_{\mathcal{C}})$, runs $\mathcal{C}$ on $(x_{\mathcal{C}}, w_{\mathcal{C}})$, and outputs $\widehat{b}_{\mathcal{C}} + \zeta + \Delta$. Since $b_{\mathcal{B}} = e + \zeta$, $b_{\mathcal{C}} = e + \zeta + \Delta$, and $e = \langle u, t\rangle + \langle v, s\rangle$, the winning event in Hyb_8 is identical to the winning event of $\widehat{\mathcal{A}_{\mathsf{moe}}}$ in the computational Decisional-Coset-MOE experiment. By Theorem 21, $\Pr[\mathsf{Hyb}_8 = 1] \leq \frac{1}{2} + \mathsf{negl}(\lambda)$. Combining the preceding hybrid transitions and the statistical error ε proves the theorem.

Changes for the identical-uniform challenge distribution. For the ID_{U}-generalized UPO special case, we replace the injective/lossy-function and universal-hashing hybrids Hyb_6–Hyb_8, together with Lemma 14, by the NCE-based hybrids of [ABH+26, Theorem 5.2]. In particular, for the common uniform challenge point x, we sample one key-robust NCE key, encrypt $\mathsf{msg}(u, v)$ under this key, and use the resulting NCE ciphertext and masked key string in both modified programs. NCE correctness and key robustness give the required functionality equivalence to invoke iO security, while the non-committing property gives a simulation of the pre-split information that is independent of (u, v). The preceding puncturable PRF and $i\mathcal{O}$ hybrids and the final reduction to computational Decisional-Coset-MOE remain unchanged. Since the required NCE can be instantiated from post-quantum one-way functions, the LWE assumption can be replaced with one-way functions in this case. For completeness, we give the formal proof of the "Moreover" part in Section 9.2.1.

□

9.2.1 Formal proof of the moreover part of Theorem 23

We need the following primitive.

Definition 41 (Key-robust one-time non-committing encryption)**.** *A key-robust one-time non-committing encryption scheme* $\mathsf{NCE} = (\mathsf{NCE.Gen}, \mathsf{NCE.Enc}, \mathsf{NCE.Dec}, \mathsf{NCE.Fake}, \mathsf{NCE.Open})$ *with message space* $\{0,1\}^{\ell_{\mathsf{NCE-msg}}}$ *and key space* $\{0,1\}^{\ell_{\mathsf{NCE-key}}}$ *consists of five PPT algorithms as follows.*

1. $k \leftarrow \mathsf{NCE.Gen}(1^\lambda)$: *outputs an encryption key* k*.*
2. $\mathsf{ct} \leftarrow \mathsf{NCE.Enc}(k, M)$: *takes as input a key* k *and a message* M*, and outputs a ciphertext* ct*.*
3. $M' \leftarrow \mathsf{NCE.Dec}(k, \mathsf{ct})$: *is deterministic and outputs a message* M' *or* $\perp$*.*
4. $(\mathsf{ct}^\star, \mathsf{st}) \leftarrow \mathsf{NCE.Fake}(1^\lambda)$: *outputs a fake ciphertext and state.*
5. $k^\star \leftarrow \mathsf{NCE.Open}(\mathsf{st}, M)$: *takes as input a state and a message, and outputs a key.*

A key k *is good if, for every message* M *and every* $k' \neq k$*,*

$$\mathsf{NCE.Dec}(k, \mathsf{NCE.Enc}(k, M)) = M \qquad \text{and} \qquad \mathsf{NCE.Dec}(k', \mathsf{NCE.Enc}(k, M)) = \perp. \tag{111}$$

A key generated by NCE.Gen *is good with overwhelming probability. Non-committing security requires that, for every QPT algorithm that outputs a message* M *together with quantum auxiliary information* R*,*

$$(R, \mathsf{ct}, k) \approx_c (R, \mathsf{ct}^\star, k^\star), \tag{112}$$

where $k \leftarrow \mathsf{NCE.Gen}(1^\lambda)$ *and* $\mathsf{ct} \leftarrow \mathsf{NCE.Enc}(k, M)$ *in the first distribution, while* $(\mathsf{ct}^\star, \mathsf{st}) \leftarrow \mathsf{NCE.Fake}(1^\lambda)$ *and* $k^\star \leftarrow \mathsf{NCE.Open}(\mathsf{st}, M)$ *in the second distribution. We may assume that* NCE.Open *is deterministic by including its random tape in* st*.*

Theorem 24 (Key-robust one-time non-committing encryption from $i\mathcal{O}$ and one-way functions [ABH+26])**.** *Assume the existence of post-quantum polynomially secure $i\mathcal{O}$ and post-quantum one-way functions. For every polynomially bounded input length $\ell_{\mathsf{in}} = \ell_{\mathsf{in}}(\lambda) \in \omega(\log \lambda)$, the parameters can be chosen so that there exist $n = n(\lambda) \in \omega(\log \lambda)$ and a key-robust one-time non-committing encryption scheme with message space $\{0,1\}^{8n}$ and key space $\{0,1\}^{\ell_{\mathsf{NCE-key}}}$, where $\ell_{\mathsf{NCE-key}} \leq \ell_{\mathsf{in}}(\lambda)$, satisfying Equations (111) and (112) against QPT distinguishers with quantum auxiliary information.*

For the proof below, fix n and a key-robust one-time non-committing encryption scheme as guaranteed by Theorem 24.

Detailed proof of the moreover part of Theorem 23. Fix a QPT triplet of adversaries against ID_{U}-generalized UPO. The challenger samples $x \xleftarrow{\$} \mathcal{X}_\lambda$ and $b \xleftarrow{\$} \{0,1\}$, gives the same point x to both recipients after the split, and uses the generalized punctured circuit that equals $\mu_{\mathcal{B}}(x)$ at x if $b = 1$ and equals $C_k(x)$ otherwise. Choose $n \in \omega(\log \lambda)$ and a key-robust one-time non-committing encryption scheme with message length $m = 8n$ and key length $\ell_{\mathsf{NCE-key}} \leq \ell_{\mathsf{in}}$. Parse every point as $x = x_0 \| x_1$, where $|x_0| = \ell_{\mathsf{NCE-key}}$.

Specialize Hyb_0–Hyb_5 above to $x_{\mathcal{B}} = x_{\mathcal{C}} = x$, $b_{\mathcal{B}} = b_{\mathcal{C}} = b$, $\Delta = 0$, and $\mu_{\mathcal{B}}$. Thus, after sampling $u \xleftarrow{\$} A \setminus \{0\}$, $v \xleftarrow{\$} A^{\perp} \setminus \{0\}$, and $\zeta \xleftarrow{\$} \{0,1\}$, we have $b = \alpha + \beta + \zeta$, where $\alpha = \langle v, s \rangle$ and $\beta = \langle u, t \rangle$. The puncturable PRF is punctured at the single point x, its value at x is replaced by a uniform y, and we sample $q \xleftarrow{\$} \mathcal{Y}_\lambda$ and write $y = q \oplus \beta \cdot \mathsf{diff}(x)$, where $\mathsf{diff}(z) := C_k(z) \oplus \mu_{\mathcal{B}}(z)$. The two hardcoded values at x are therefore $q \oplus C_k(x) \oplus (\alpha + \zeta) \cdot \mathsf{diff}(x)$ and $q \oplus \beta \cdot \mathsf{diff}(x)$.

Next sample $k_{\mathsf{nce}} \leftarrow \mathsf{NCE.Gen}(1^\lambda)$, $\mathsf{ct} \leftarrow \mathsf{NCE.Enc}(k_{\mathsf{nce}}, \mathsf{msg}(u,v))$, and set $\mathsf{str} := k_{\mathsf{nce}} \oplus x_0$. Replace the two point-hardcoded programs by the following programs, initially using the punctured key $K_{\{x\}}$ in their last lines.

$\underline{S_0(a,z)}$

Hardcoded: $\hat{M}_0, C_k, \mu_{\mathcal{B}}, K_{\{x\}}, q, \zeta, \mathsf{ct}, \mathsf{str}, x_1$

1. If $\hat{M}_0(a) = 0$, output $\perp$.
2. Parse $z = z_0 \| z_1$, set $k_z := \mathsf{str} \oplus z_0$, and compute $M_z := \mathsf{NCE.Dec}(k_z, \mathsf{ct})$.
3. If $z_1 = x_1$ and $M_z \neq \perp$, parse $(u', v') := M_z$ and output $q \oplus C_k(z) \oplus (\langle v', a \rangle + \zeta) \cdot \mathsf{diff}(z)$.
4. Output $F_{K_{\{x\}}}(z) \oplus C_k(z)$.

$\underline{S_1(a,z)}$

Hardcoded: $\hat{M}_1, C_k, \mu_{\mathcal{B}}, K_{\{x\}}, q, \mathsf{ct}, \mathsf{str}, x_1$

1. If $\hat{M}_1(a) = 0$, output $\perp$.
2. Parse $z = z_0 \| z_1$, set $k_z := \mathsf{str} \oplus z_0$, and compute $M_z := \mathsf{NCE.Dec}(k_z, \mathsf{ct})$.
3. If $z_1 = x_1$ and $M_z \neq \perp$, parse $(u', v') := M_z$ and output $q \oplus \langle u', a \rangle \cdot \mathsf{diff}(z)$.
4. Output $F_{K_{\{x\}}}(z)$.

Condition on the overwhelmingly likely event that k_{nce} is good. At $z = x$, $k_z = \mathsf{str} \oplus x_0 = k_{\mathsf{nce}}$, so correctness gives $M_z = \mathsf{msg}(u,v)$. Hence, for $a \in A + s$, S_0 returns $q \oplus C_k(x) \oplus (\alpha + \zeta) \cdot \mathsf{diff}(x)$,

and for $a \in A^{\perp} + t$, S_1 returns $q \oplus \beta \cdot \mathsf{diff}(x)$. If $z \neq x$, either $z_1 \neq x_1$, or $z_0 \neq x_0$ and $k_z \neq k_{\mathsf{nce}}$, in which case key robustness gives $M_z = \perp$. Thus the relevant circuits have the same functionality, and two applications of $i\mathcal{O}$-security justify the replacement. We then replace $K_{\{x\}}$ by the full puncturable PRF key K in the last line of each program. The relevant circuits again have the same functionality, and this replacement follows from $i\mathcal{O}$-security.

Since x_0 is uniform, the same distribution can be sampled by first generating k_{nce} and ct, then sampling $\mathsf{str} \xleftarrow{\$} \{0,1\}^{\ell_{\mathsf{NCE-key}}}$ and $x_1 \xleftarrow{\$} \{0,1\}^{\ell_{\mathsf{in}} - \ell_{\mathsf{NCE-key}}}$, and finally setting $x_0 := \mathsf{str} \oplus k_{\mathsf{nce}}$. Apply non-committing security to replace $(\mathsf{ct}, k_{\mathsf{nce}})$ by $(\mathsf{ct}^\star, k^\star)$, where $(\mathsf{ct}^\star, \mathsf{st}) \leftarrow \mathsf{NCE.Fake}(1^\lambda)$ and $k^\star := \mathsf{NCE.Open}(\mathsf{st}, \mathsf{msg}(u,v))$. Except with negligible probability, $\mathsf{NCE.Dec}(k^\star, \mathsf{ct}^\star) = \mathsf{msg}(u,v)$; otherwise, testing this equality would distinguish the fake/open distribution from the real correct distribution. In the resulting hybrid, the values $(\mathsf{ct}^\star, \mathsf{str}, x_1)$ given before the split are independent of (u,v).

We reduce this hybrid to the computational Decisional-Coset-MOE experiment without auxiliary information. The algorithm $\widehat{\mathcal{A}}$ receives $|A_{s,t}\rangle$, $\hat{M}_0, \hat{M}_1$, runs the first stage of $\mathcal{A}$ to obtain $(k, \mu_{\mathcal{B}}, \mu_{\mathcal{C}})$, samples $(\mathsf{ct}^\star, \mathsf{st}) \leftarrow \mathsf{NCE.Fake}(1^\lambda)$, $\mathsf{str} \xleftarrow{\$} \{0,1\}^{\ell_{\mathsf{NCE-key}}}$, $x_1 \xleftarrow{\$} \{0,1\}^{\ell_{\mathsf{in}} - \ell_{\mathsf{NCE-key}}}$, $q \xleftarrow{\$} \mathcal{Y}_\lambda$, $\zeta \xleftarrow{\$} \{0,1\}$, and $K \leftarrow \mathsf{PRF.Gen}(1^\lambda)$. It constructs and obfuscates the two programs S_0, S_1 with $\mathsf{ct}^\star$ and the full key K, gives the resulting UPO state to the continuation of $\mathcal{A}$, and appends classical copies of $(\mathsf{st}, \mathsf{str}, x_1, \zeta)$ to both output registers.

After the split, each of $\widehat{\mathcal{B}}, \widehat{\mathcal{C}}$ receives (u,v). It computes $M := \mathsf{msg}(u,v)$, $k^\star := \mathsf{NCE.Open}(\mathsf{st}, M)$, $x_0 := \mathsf{str} \oplus k^\star$, and $x := x_0 \| x_1$, runs the corresponding original recipient on x, and adds ζ to its output. Since the IDU challenge bit is $b = \langle u,t \rangle + \langle v,s \rangle + \zeta$, both original recipients output b if and only if both $\widehat{\mathcal{B}}$ and $\widehat{\mathcal{C}}$ output $\langle u,t \rangle + \langle v,s \rangle$. Therefore, Corollary 3 bounds the winning probability by $\frac{1}{2} + \mathsf{negl}(\lambda)$. Combining the puncturable PRF, $i\mathcal{O}$, non-committing encryption, and negligible bad-key errors proves $\mathsf{ID_U}$-generalized UPO security from post-quantum $i\mathcal{O}$ and post-quantum one-way functions. □

10 Applications of Correlated UPO to Copy-Protection

We recall the UPO-to-copy-protection implications of [AB24, ABH+26] and adapt them to correlated challenge distributions. Throughout this subsection, $\mathsf{Circ} = \{C_k : k \in \mathcal{K}_\lambda\}$ is a keyed circuit class and $\mathcal{D}_{\mathsf{Circ}}$ is induced by an efficiently samplable key distribution $\mathcal{D}_K$: sampling $C \leftarrow \mathcal{D}_{\mathsf{Circ}}(1^\lambda)$ means sampling $k \leftarrow \mathcal{D}_K(1^\lambda)$ and setting $C := C_k$. The key, and hence C, is sampled independently of the corresponding challenge sampler.

Definition 42 (Two-point unpredictability-style puncturing security, adapted from [ABH+26])**.** *Let* Circ *be a circuit class with output space* $\mathcal{Y}_\lambda = \{0,1\}^{m(\lambda)}$*. The class* Circ *is two-point unpredictability-style puncturable with respect to* $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{CP}})$ *if there exists a deterministic PPT algorithm* $\mathsf{Puncture}$ *such that, on input a circuit* $C \in \mathsf{Circ}_\lambda$ *and a set* $\mathcal{S} \subseteq \mathcal{X}_\lambda$ *of size at most two, it outputs a circuit* $C^{\setminus \mathcal{S}} : \mathcal{X}_\lambda \to \mathcal{Y}_\lambda$ *satisfying*

$$C^{\setminus \mathcal{S}}(x) = \begin{cases} 0^{m(\lambda)}, & x \in \mathcal{S}, \\ C(x), & x \notin \mathcal{S}. \end{cases} \tag{113}$$

All circuits compared below are padded to the same size. Moreover, for every QPT adversary $\mathcal{D}$*, there exists a negligible function* negl *such that*

$$\Pr\left[\widehat{y}_{\mathcal{B}} = C(x_{\mathcal{B}}) \ \vee \ \widehat{y}_{\mathcal{C}} = C(x_{\mathcal{C}}) : \begin{array}{l} C \leftarrow \mathcal{D}_{\mathsf{Circ}}(1^\lambda), \\ (x_{\mathcal{B}}, x_{\mathcal{C}}) \leftarrow \mathsf{Samp}_{\mathsf{CP}}(1^\lambda), \\ \mathcal{S} := \{x_{\mathcal{B}}, x_{\mathcal{C}}\}, \\ C^{\setminus \mathcal{S}} \leftarrow \mathsf{Puncture}(C, \mathcal{S}), \\ (\widehat{y}_{\mathcal{B}}, \widehat{y}_{\mathcal{C}}) \leftarrow \mathcal{D}(C^{\setminus \mathcal{S}}, x_{\mathcal{B}}, x_{\mathcal{C}}) \end{array}\right] \leq 1 - \left(1 - 2^{-m(\lambda)}\right)^2 + \mathsf{negl}(\lambda). \tag{114}$$

Definition 43 (Two-point pseudorandomness-style puncturing security with correlated challenges, adapted from [ABH+26])**.** *Let* Circ *be a deterministic circuit class with output space* $\mathcal{Y}_\lambda = \{0,1\}^{m(\lambda)}$*. We say that* Circ *is two-point pseudorandomness-style puncturable with respect to* $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{PRCP}})$ *if it has a deterministic PPT puncturing algorithm satisfying Equation (113), and the following two ensembles are computationally indistinguishable.*

Sample $C \leftarrow \mathcal{D}_{\mathsf{Circ}}(1^\lambda)$ *and* $(x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}) \leftarrow \mathsf{Samp}_{\mathsf{PRCP}}(1^\lambda)$*, and let* $\mathcal{S} := \{x_{\mathcal{B}}, x_{\mathcal{C}}\}$*. In the first ensemble, output*

$$(C^{\setminus \mathcal{S}}, x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}, C(x_{\mathcal{B}}), C(x_{\mathcal{C}})). \tag{115}$$

In the second ensemble, sample an independent $U_x \xleftarrow{\$} \mathcal{Y}_\lambda$ *for every distinct* $x \in \mathcal{S}$*, and output*

$$(C^{\setminus \mathcal{S}}, x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}, U_{x_{\mathcal{B}}}, U_{x_{\mathcal{C}}}). \tag{116}$$

Given a UPO scheme $\Pi = (\mathsf{Obf}, \mathsf{Eval})$, define CP_Π by $\mathsf{CopyProtect}(1^\lambda, C) := \mathsf{Obf}(1^\lambda, C)$ and by letting $\mathsf{CP}_\Pi.\mathsf{Eval}(\rho, x)$ run $(\rho', y) \leftarrow \mathsf{Eval}(\rho, x)$ and output y. Whenever we invoke UPO security below, the supported circuit class is understood to contain the base circuits and the circuits produced by the relevant puncturing and programming algorithms, padded to a common size. For Definition 42, the corresponding generalized-puncturing algorithm first computes $C^{\setminus \mathcal{S}} \leftarrow \mathsf{Puncture}(C, \mathcal{S})$ and then programs $\mu_{\mathcal{B}}(x_{\mathcal{B}})$ and, if $x_{\mathcal{C}} \neq x_{\mathcal{B}}$, $\mu_{\mathcal{C}}(x_{\mathcal{C}})$, using Bob priority on a collision. Thus the punctured event of the UPO game uses the same puncturing procedure as the security assumption.

We first adapt the non-oracular point-function application of generalized UPO to arbitrary high-min-entropy point distributions.

Theorem 25 (Copy protection for point functions from $\mathsf{ID}_{\mathcal{D}}$-generalized UPO, adapted from [AB24, Theorem 89 and Corollary 90])**.** *Let $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda\in\mathbb{N}}$ be a QPT-samplable distribution over $\mathcal{X}_\lambda = \{0,1\}^{\ell_{\mathsf{in}}(\lambda)}$ such that $\widetilde{H}_\infty(\mathcal{D}_\lambda) \in \omega(\log\lambda)$. Suppose that there is a UPO scheme for polynomial-size circuits satisfying $\mathsf{ID}_{\mathcal{D}}$-generalized UPO security and $i\mathcal{O}$-security. Then point functions whose point is sampled according to $\mathcal{D}_\lambda$ have a copy-protection scheme satisfying identical-challenge point-function security with respect to $\mathcal{D}$ from Definition 32.*

Proof. We adapt the proof of [AB24, Theorem 89]. Let C_0 be the all-zero circuit and let G_y be obtained by programming C_0 to output one at y. The $i\mathcal{O}$-security of UPO replaces an obfuscation of P_y by an obfuscation of the functionally equivalent circuit G_y. In the independent event, replace the desired output by zero. This changes success only when the independent samples $x, y \leftarrow \mathcal{D}_\lambda$ coincide, which occurs with probability

$$\mathrm{cp}(\mathcal{D}_\lambda) := \sum_x \Pr[\mathcal{D}_\lambda = x]^2 \le 2^{-\widetilde{H}_\infty(\mathcal{D}_\lambda)} = \mathsf{negl}(\lambda). \tag{117}$$

The switch from an obfuscation of G_y to an obfuscation of C_0 in the independent event is precisely the guarantee given by $\mathsf{ID}_{\mathcal{D}}$-generalized UPO security with the constant-one replacement circuits. After this switch, both events have challenge-point marginal $\mathcal{D}_\lambda$, and the final experiment is the $\mathsf{ID}_{\mathcal{D}}$-generalized UPO experiment. Thus every adversary succeeds with probability at most $1/2+\mathsf{negl}(\lambda)$. By Lemma 7, this implies the required trivial-adversary comparison. The uniform specialization is exactly [AB24, Corollary 90]. □

Corollary 4 (Copy-protection for point functions)**.** *Assume post-quantum polynomially secure $i\mathcal{O}$ and the quantum hardness of LWE. For every constant $c > 0$ and every QPT-samplable distribution ensemble $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda\in\mathbb{N}}$ over $\mathcal{X}_\lambda$ satisfying $\widetilde{H}_\infty(\mathcal{D}_\lambda) \ge \lambda^c$, point functions whose point is sampled according to $\mathcal{D}_\lambda$ have a copy-protection scheme satisfying identical-challenge point-function security with respect to $\mathcal{D}$. Moreover, when $\mathcal{D}_\lambda$ is uniform over $\{0,1\}^{\ell_{\mathsf{in}}(\lambda)}$ and $\ell_{\mathsf{in}} \in \omega(\log\lambda)$, LWE can be replaced with post-quantum one-way functions.*

Proof. The first claim follows from Theorems 18 and 23 and Lemma 6, which give a UPO satisfying $\mathsf{ID}_{\mathcal{D}}$-generalized UPO security and $i\mathcal{O}$-security, and Theorem 25. The uniform claim uses the "Moreover" part of Theorem 23 together with the add-$i\mathcal{O}$ transformation of [ABH+26, Lemma 3.16]. □

Corollary 5 (Copy protection for k-point functions)**.** *Assume post-quantum polynomially secure $i\mathcal{O}$ and the quantum hardness of LWE. Let $k = k(\lambda)$ be polynomially bounded, and let $\mathcal{D}^{(k)} = \{\mathcal{D}^{(k)}_\lambda\}_{\lambda\in\mathbb{N}}$ be a QPT-samplable distribution over the k-element subsets of $\mathcal{X}_\lambda$ such that $\widetilde{H}_\infty(\mathcal{D}^{(k)}_\lambda) \ge \lambda^c$ for some constant $c > 0$. Let $\mathcal{D}^{\mathsf{pt}}_\lambda$ be the distribution obtained by sampling $S \leftarrow \mathcal{D}^{(k)}_\lambda$ and then $x \xleftarrow{\$} S$, and suppose that $\widetilde{H}_\infty(\mathcal{D}^{\mathsf{pt}}_\lambda) \ge \lambda^{c'}$ for some constant $c' > 0$. Suppose moreover that the resulting distribution over k-point functions is preimage-samplable in the sense of [AB24, Definition 84], with the uniform challenge distribution in that definition replaced by $\mathcal{D}^{\mathsf{pt}}_\lambda$. Then the k-point functions sampled according to $\mathcal{D}^{(k)}$ have a copy-protection scheme satisfying the identical-challenge security (see Remark 4). If $\mathcal{D}^{(k)}_\lambda$ is uniform over the k-element subsets of $\mathcal{X}_\lambda$ and $k(\lambda)/|\mathcal{X}_\lambda|$ is negligible, then LWE can be replaced with post-quantum one-way functions.*

Proof. For every fixed $x \in \mathcal{X}_\lambda$,

$$\Pr_{S\leftarrow\mathcal{D}^{(k)}_\lambda}[x \in S] = k(\lambda)\Pr\left[\mathcal{D}^{\mathsf{pt}}_\lambda = x\right] \le k(\lambda)2^{-\widetilde{H}_\infty(\mathcal{D}^{\mathsf{pt}}_\lambda)} = \mathsf{negl}(\lambda), \tag{118}$$

so the sampled function class is evasive. Under the stated preimage-samplability assumption, the proof of [AB24, Theorem 89] applies after replacing each uniform challenge point by a sample from $\mathcal{D}_\lambda^{\mathsf{pt}}$. By Theorems 18 and 23 and Lemma 6, the stated assumptions give the required $\mathsf{ID}_{\mathcal{D}^{\mathsf{pt}}}$-generalized UPO and $i\mathcal{O}$-security. The uniform specialization is the setting of [AB24, Theorem 85 and Corollary 90], and the "Moreover" part of Theorem 23 together with the add-$i\mathcal{O}$ transformation of [ABH+26, Lemma 3.16] supplies the required ID_{U}-generalized UPO from one-way functions. The trivial success probability is given by Equation (64). □

Theorem 26 (Copy protection for compute-and-compare functions)**.** *Assume post-quantum polynomially secure $i\mathcal{O}$ and the quantum hardness of LWE. Let* Samp *be a QPT algorithm that samples $(f, y) \leftarrow \mathsf{Samp}(1^\lambda)$, where $f : \mathcal{X}_\lambda \to \mathcal{Z}_\lambda$ is a polynomial-size circuit, $\mathcal{Z}_\lambda = \{0,1\}^{\ell_{\mathsf{cc}}(\lambda)}$, and $y \in \mathcal{Z}_\lambda$. Let (F, Y) denote its output distribution and suppose that $\widetilde{H}_\infty(Y \mid F) \geq \lambda^c$ for some constant $c > 0$. Fix a family of distributions $\{\mathcal{Q}_{f,z}\}_{f,z}$ as in Definition 33, and suppose that the resulting identical-challenge experiment is efficiently realizable. Then the compute-and-compare functions sampled by* Samp *have a copy-protection scheme satisfying the identical-challenge security of Definition 33. Moreover, if Y conditioned on $F = f$ is uniform over $\mathcal{Z}_\lambda$ for every f in the support of F, then the LWE assumption can be replaced with post-quantum one-way functions.*

Proof. Let Π be a UPO, guaranteed by Theorems 18 and 23 for any polynomial input and output length circuit class, and let $\mathsf{CP}_{\mathsf{PF}}$ be the point-function scheme obtained from Π by the construction used in Theorem 25. Following the generic transformation of [CMP24, Construction 4 and Theorem 7], define

$$\rho_y \leftarrow \mathsf{CopyProtect}_{\mathsf{PF}}(1^\lambda, P_y), \qquad \mathsf{CopyProtect}_{\mathsf{CC}}(1^\lambda, \mathsf{CC}_{f,y}) := (f, \rho_y), \tag{119}$$

and

$$\mathsf{Eval}_{\mathsf{CC}}((f, \rho_y), x) := \mathsf{Eval}_{\mathsf{PF}}(\rho_y, f(x)). \tag{120}$$

Correctness follows from $\mathsf{CC}_{f,y}(x) = P_y(f(x))$.

For every f in the support of F, let $\mathcal{D}_f$ denote the conditional distribution of Y given $F = f$. Consider the correlated-UPO sampler that samples

$$(f, z) \leftarrow \mathsf{Samp}(1^\lambda), \qquad x \leftarrow \mathcal{Q}_{f,z}, \qquad b \xleftarrow{\$} \{0,1\}, \tag{121}$$

where the second output of Samp is denoted by z, and outputs

$$(z_{\mathsf{aux}}, x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}, w_{\mathcal{B}}, w_{\mathcal{C}}) := (f, z, z, b, b, x, x). \tag{122}$$

This sampler satisfies Definition 26: the shared bit is uniform conditioned on (f, z, x) and

$$\widetilde{H}_\infty(z \mid f) = \widetilde{H}_\infty(Y \mid F) \geq \lambda^c.$$

We adapt the proof of Theorem 25. Let C_0 be the all-zero circuit and let G_y be obtained by programming C_0 to output one at y. The $i\mathcal{O}$-security of UPO replaces an obfuscation of P_y by an obfuscation of the functionally equivalent circuit G_y. In the event $b = 0$, replace the desired output by zero. This changes success only if the independent samples $y, z \leftarrow \mathcal{D}_f$ coincide. Averaged over f, this occurs with probability

$$\mathbb{E}_f[\mathrm{cp}(\mathcal{D}_f)] \leq 2^{-\widetilde{H}_\infty(Y|F)} = \mathsf{negl}(\lambda). \tag{123}$$

Next, in the event $b = 0$, replace the obfuscation of G_y by an obfuscation of C_0. To justify this change, consider the correlated-UPO sampler that samples $(f, y) \leftarrow \mathsf{Samp}(1^\lambda)$, a uniform bit $d \xleftarrow{\$} \{0,1\}$, and outputs

$$(z_{\mathsf{aux}}, x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}, w_{\mathcal{B}}, w_{\mathcal{C}}) := (f, y, y, d, d, \bot, \bot). \tag{124}$$

This is a valid sampler because $\widetilde{H}_\infty(y \mid f) \geq \lambda^c$. Correlated UPO security therefore implies that the ensembles consisting of $(f, \mathsf{Obf}(G_y))$ and $(f, \mathsf{Obf}(C_0))$ are computationally indistinguishable: otherwise, $\mathcal{R}_{\mathcal{A}}$ runs the distinguisher and sends a classical copy of its output to each of $\mathcal{R}_{\mathcal{B}}$ and $\mathcal{R}_{\mathcal{C}}$. In the event $b = 0$, the fresh value $z \leftarrow \mathcal{D}_f$ and the point $x \leftarrow \mathcal{Q}_{f,z}$ are independent of y conditioned on f and are efficiently sampleable. Hence adjoining (z, x) and applying the QPT algorithms of the compute-and-compare adversary preserves this indistinguishability.

After this change, the resulting experiment is exactly the correlated-UPO experiment for the first sampler (Equation (122)) displayed above. The reduction $\mathcal{R}_{\mathcal{A}}$, after receiving f and the UPO state, gives (f, ρ) to $\mathcal{A}$ and forwards its two output registers; after the split, $\mathcal{R}_{\mathcal{B}}$ and $\mathcal{R}_{\mathcal{C}}$ ignore the common UPO point z and give the common post-split auxiliary string x to $\mathcal{B}$ and $\mathcal{C}$, respectively. When the shared bit is zero, the state protects C_0 and both recipients must output zero. When the shared bit is one, the common UPO challenge point is z, the state protects $G_z = P_z$, and both recipients must output one. Since $x \leftarrow \mathcal{Q}_{f,z}$ and hence $f(x) = z$, their target value is

$$\mathsf{CC}_{f,y}(x) = P_y(z). \tag{125}$$

Thus correlated UPO security implies that every adversary in the last experiment succeeds with probability at most $1/2 + \mathsf{negl}(\lambda)$. Therefore, we conclude that the success probability of the adversaries in the original compute-and-compare copy protection experiment is at most $1/2 + \mathsf{negl}'(\lambda)$ for some negligible function $\mathsf{negl}'(\lambda)$.

Combining the last conclusion with the fact that the trivial adversarial strategy of outputting 1 (i.e., $\mathcal{A}$ sends the copy-protected state to $\mathcal{B}$ and $\mathcal{C}$ always outputs 1) has success probability

$$\frac{1}{2} + \frac{1}{2}\mathbb{E}_f[\mathrm{cp}(\mathcal{D}_f)] \leq \frac{1}{2} + \frac{1}{2}\,2^{-\widetilde{H}_\infty(Y|F)}, \tag{126}$$

which is $1/2$ plus some negligible function, we conclude that the trivial success probability for the experiment is $\frac{1}{2}$ plus some negligible function. Hence, by the preceding argument, we conclude that the success probability for any adversary is upper bounded by the trivial success probability upto negligible corrections (see Definition 33), which concludes the proof of the security.

The "Moreover" clearly follows by noting that in this special case, we only need correlated UPO secure against identical uniform challenge distribution and then replacing the general statement of Theorem 23 in the above proof with the "Moreover" part of Theorem 23. □

Next, we turn our attention to unpredictability-style copy protection for puncturable functionalities.

Theorem 27 (Unpredictability-style copy protection from UPO, adapted from [AB24, Theorem 57])**.** *Let $\mathcal{D}_X$ be the joint distribution induced by $\mathsf{Samp}_{\mathsf{CP}}$, let $m \in \omega(\log \lambda)$, and let Circ be two-point unpredictability-style puncturable with respect to $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{CP}})$. If Π satisfies $\mathcal{D}_X$-generalized UPO security, then CP_Π satisfies unpredictability-style copy-protection security with respect to $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{CP}})$ (see Definition 31).*

Proof. We adapt the reduction of [AB24, Theorem 57] while preserving the sampled joint pair. In summary, the argument below uses the challenge pair as a jointly sampled object, and hence permits arbitrary correlation among challenge points.

Fix a copy-protection adversary, and construct a generalized-UPO adversary using the constant replacement circuits $\mu_X(z) := 0^{m(\lambda)}$ for $X \in \{\mathcal{B}, \mathcal{C}\}$. Let p_0 be the probability that both recipients recover their respective values when the UPO challenge bit is zero, and let p_1 be the probability that neither recipient recovers its respective value when the challenge bit is one. By construction, p_0 is exactly the success probability in the copy-protection experiment (see Definition 31). In the punctured event, the puncturing guarantee Equation (114) yields

$$p_1 \geq \left(1 - 2^{-m(\lambda)}\right)^2 - \mathsf{negl}(\lambda). \tag{127}$$

The reduction outputs zero precisely upon a correct prediction and one otherwise, so generalized-UPO security gives $(p_0 + p_1)/2 \leq 1/2 + \mathsf{negl}(\lambda)$. Therefore,

$$p_0 \leq 1 - \left(1 - 2^{-m(\lambda)}\right)^2 + \mathsf{negl}(\lambda), \tag{128}$$

which is negligible in λ. □

Corollary 6 (Unpredictability-style copy protection with correlated challenges)**.** *Assume post-quantum polynomially secure $i\mathcal{O}$ and the quantum hardness of LWE. Let $c > 0$, let $\mathsf{Samp}_{\mathsf{CP}}$ be any QPT sampler whose two challenge-point marginals have min-entropy at least λ^c, and let Circ have output length $m \in \omega(\log \lambda)$ and satisfy Definition 42. Then there exists a copy-protection scheme for Circ satisfying unpredictability-style security for $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{CP}})$.*

Proof. Apply Theorem 23 and Lemma 6 to the joint distribution induced by $\mathsf{Samp}_{\mathsf{CP}}$, and then apply Theorem 27. □

We note that the above result can be strengthened to oracular security when the two challenge points are identical and are sampled from a high-min-entropy distribution. Let $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ be a QPT-samplable distribution ensemble over $\mathcal{X}_\lambda$. We write $\mathsf{Samp}_{\mathsf{ID}_\mathcal{D}}$ for the sampler that samples $x \leftarrow \mathcal{D}_\lambda$ and outputs $(x_\mathcal{B}, x_\mathcal{C}) := (x, x)$.

Theorem 28 (Oracular unpredictability-style copy protection for identical challenges)**.** *Let $\widetilde{H}_\infty(\mathcal{D}_\lambda) \in \omega(\log \lambda)$ and $m \in \omega(\log \lambda)$. Suppose that Circ is two-point unpredictability-style puncturable with respect to $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{ID}_\mathcal{D}})$ and that Π satisfies $\mathsf{ID}_\mathcal{D}$-generalized UPO security. Then CP_Π satisfies oracular unpredictability-style copy-protection security with respect to $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{ID}_\mathcal{D}})$.*

Proof. We combine the reduction above with the oracle-hybrid argument of [ABH+26, Remark 8.3], specialized to one common point $x \leftarrow \mathcal{D}_\lambda$. Let G_x be the punctured circuit, which agrees with C off x and outputs $0^{m(\lambda)}$ at x. Conditioned on the event that the UPO challenge bit is one, suppose that $\mathcal{A}$ makes at most $q = q(\lambda)$ oracle queries, and let W be their total query weight on x when $\mathcal{A}$ receives $\mathsf{Obf}(G_x)$ but still has oracle access to C. We claim that $\mathbb{E}[W]$ is negligible.

To see this, let $\delta_\lambda := 2^{-\widetilde{H}_\infty(\mathcal{D}_\lambda)}$. For $q \geq 1$, construct an $\mathsf{ID}_\mathcal{D}$-generalized UPO adversary $\mathcal{R} = (\mathcal{R}_\mathcal{A}, \mathcal{R}_\mathcal{B}, \mathcal{R}_\mathcal{C})$ that samples $k \leftarrow \mathcal{D}_K(1^\lambda)$, submits k together with the constant replacement circuits $\mu_\mathcal{B}(z) = \mu_\mathcal{C}(z) := 0^{m(\lambda)}$, and uses $C := C_k$ to simulate $\mathcal{A}$'s C-oracle. The algorithm $\mathcal{R}_\mathcal{A}$ chooses $i \xleftarrow{\$} [q]$, measures the input of the i-th query if it occurs, sets $x' := \perp$ otherwise, and sends x' to both reduction recipients. Upon receiving the common challenge x, each of $\mathcal{R}_\mathcal{B}$ and $\mathcal{R}_\mathcal{C}$ outputs one if $x' = x$ and zero otherwise. Conditioned on the event that the UPO challenge

bit is zero, x' is independent of the fresh $x \leftarrow \mathcal{D}_\lambda$, so both reduction recipients are correct with probability at least $1 - \delta_\lambda$. Conditioned on the event that the UPO challenge bit is one, both are correct with probability $\mathbb{E}[W]/q$. Therefore, $\mathsf{ID}_{\mathcal{D}}$-generalized UPO security gives

$$\frac{1}{2}\left(1 - \delta_\lambda + \frac{\mathbb{E}[W]}{q}\right) \leq \frac{1}{2} + \mathsf{negl}(\lambda), \tag{129}$$

and hence $\mathbb{E}[W] \leq q\delta_\lambda + \mathsf{negl}(\lambda) = \mathsf{negl}(\lambda)$. The case $q = 0$ is immediate.

By the BBBV query bound [ABH+26, Theorem 2.1], changing $\mathcal{A}$'s oracle from C to G_x changes the adversarial success probability by at most $O(\sqrt{q\mathbb{E}[W]}) = \mathsf{negl}(\lambda)$. Both post-split recipient oracles are already identical, since $C^{\setminus x} = G_x^{\setminus x}$. The puncturing reduction can consequently simulate the punctured event using G_x, and the same p_0, p_1 calculation as above gives negligible piracy success and hence the required trivial-strategy bound. □

Corollary 7 (Oracular unpredictability-style copy protection for identical challenges)**.** *Assume post-quantum polynomially secure $i\mathcal{O}$ and the quantum hardness of LWE. Let $c > 0$ be a constant, let $\mathcal{D}$ be QPT-samplable with $\widetilde{H}_\infty(\mathcal{D}_\lambda) \geq \lambda^c$, and let* Circ *have output length $m \in \omega(\log \lambda)$ and satisfy Definition 42 with respect to $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{ID}_{\mathcal{D}}})$. Then there exists a copy-protection scheme for* Circ *satisfying oracular unpredictability-style security with respect to $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{ID}_{\mathcal{D}}})$.*

Proof. Apply Theorem 23 and Lemma 6 to $\mathsf{ID}_{\mathcal{D}}$ and then apply Theorem 28. □

For the OWF-based instantiation, we retain the identical-uniform distribution, since the "Moreover" part of Theorem 23 establishes ID_{U}-generalized UPO security.

Corollary 8 (Identical-uniform oracular unpredictability-style copy protection)**.** *Assume post-quantum polynomially secure $i\mathcal{O}$ and post-quantum one-way functions. Let $\ell_{\mathsf{in}}, m \in \omega(\log \lambda)$, and suppose that* Circ *is two-point unpredictability-style puncturable for the identical-uniform sampler. Then there exists a copy-protection scheme for* Circ *satisfying oracular unpredictability-style security for that sampler.*

Proof. Combine the "Moreover" part of Theorem 23 with Theorem 28, taking $\mathcal{D}_\lambda$ to be uniform over $\mathcal{X}_\lambda$. □

Next, we consider applications to pseudorandomness-style copy protection.

Theorem 29 (Pseudorandomness-style copy protection from correlated UPO, adapted from [ABH+26, Theorem 7.4])**.** *Let $\mathsf{Samp}_{\mathsf{PRCP}}$ be a correlated pseudorandomness-style challenge sampler, and let* Circ *satisfy two-point pseudorandomness-style puncturing security for $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{PRCP}})$. Suppose that Π satisfies correlated-UPO security for $\mathsf{Samp}_{\mathsf{PRCP}}$ and $i\mathcal{O}$-security. Then CP_Π satisfies pseudorandomness-style copy-protection security for $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{PRCP}})$.*

Proof. We adapt the non-oracular portion of [ABH+26, Theorem 7.4]. Since the challenge tuple is independent of the sampled circuit and of the adversary's pre-challenge view, we may sample it at the outset and withhold it until after the split without changing the experiment. For a sampled tuple let $\mathcal{S} := \{x_{\mathcal{B}}, x_{\mathcal{C}}\}$ contain each distinct point once. Define the effective bit $d_{x_{\mathcal{B}}} := b_{\mathcal{B}}$ and, when $x_{\mathcal{C}} \neq x_{\mathcal{B}}$, $d_{x_{\mathcal{C}}} := b_{\mathcal{C}}$, and let $T := \{x \in \mathcal{S} : d_x = 1\}$. This is precisely the Bob-priority convention on a collision. Write $R_x := C(x)$, and let $U_x, Y_x \xleftarrow{\$} \mathcal{Y}_\lambda$ be independent for each $x \in \mathcal{S}$. For values $v = (v_x)_{x \in \mathcal{S}}$, let $\mathsf{Program}(C^{\setminus \mathcal{S}}, v)$ agree with $C^{\setminus \mathcal{S}}$ off $\mathcal{S}$ and output v_x at $x \in \mathcal{S}$.

First, $i\mathcal{O}$-security replaces $\mathsf{Obf}(C)$ by $\mathsf{Obf}(\mathsf{Program}(C^{\setminus\mathcal{S}}, R))$. We use the following consequence of two-point pseudorandomness-style puncturing: for every $T \subseteq \mathcal{S}$ efficiently determined from the public tuple $(x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}})$,

$$(C^{\setminus\mathcal{S}}, R_{\mathcal{S}\setminus T}, R_T) \approx_c (C^{\setminus\mathcal{S}}, R_{\mathcal{S}\setminus T}, U_T). \tag{130}$$

Indeed, the real and fully uniform ensembles are indistinguishable by Definition 43. Applying to both the map that retains the coordinates outside T and freshly resamples those in T shows that the mixed and fully uniform ensembles are indistinguishable. The triangle inequality gives Equation (130). The sampled bits are retained throughout, so this argument also permits arbitrary correlation between the points and bits.

Apply Equation (130) to replace the protected values R_T by U_T, retaining the independent challenge values Y_T as auxiliary information. Swapping the names of the independent uniform families U_T, Y_T is distribution preserving. Apply Equation (130) a second time, now to replace the uniform challenge values by R_T while retaining the independently programmed values as auxiliary information. Thus the classical challenge is $C(x_X)$ for both recipients, while the protected circuit is programmed with independent uniform values exactly at the points whose effective bit is one. Finally, $i\mathcal{O}$-security replaces this circuit by the canonical correlated-punctured circuit.

For completeness, the final correlated-UPO reduction samples $k \leftarrow \mathcal{D}_K$ and independent $u_{\mathcal{B}}, u_{\mathcal{C}} \xleftarrow{\$} \mathcal{Y}_\lambda$ before sending $(k, \mu_{\mathcal{B}}, \mu_{\mathcal{C}})$ to the challenger, where μ_X is the constant-u_X circuit. After the split, recipient X computes $C_k(x_X)$ and feeds $(x_X, C_k(x_X))$ to its copy-protection recipient. If the points collide, $u_{\mathcal{C}}$ is unused by Bob priority. Hence the simulated experiment is exactly the final hybrid, and correlated-UPO security bounds its success by $1/2 + \mathsf{negl}(\lambda)$. No oracle-switch hybrid is used, so no independence between $x_{\mathcal{B}}$ and $x_{\mathcal{C}}$ is required. □

Corollary 9 (Pseudorandomness-style copy protection with correlated challenges)**.** *Assume post-quantum polynomially secure $i\mathcal{O}$ and the quantum hardness of LWE. Let $c > 0$, let $\mathsf{Samp}_{\mathsf{PRCP}}$ have point marginals of min-entropy at least λ^c, and suppose that Circ satisfies Definition 43. Then Circ has a copy-protection scheme satisfying pseudorandomness-style security for $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}_{\mathsf{PRCP}})$.*

Proof. Combine Theorems 18 and 23 with Theorem 29. □

For the oracular pseudorandomness-style result, let $\mathsf{Samp}^{\mathsf{pr}}_{\mathsf{ID}_{\mathcal{D}}}$ sample $x \leftarrow \mathcal{D}_\lambda$ and $b \xleftarrow{\$} \{0,1\}$ and output

$$(x_{\mathcal{B}}, x_{\mathcal{C}}, b_{\mathcal{B}}, b_{\mathcal{C}}) := (x, x, b, b). \tag{131}$$

Theorem 30 (Oracular pseudorandomness-style copy protection for identical challenges, adapted from [ABH+26, Theorem 7.4])**.** *Let $\widetilde{H}_\infty(\mathcal{D}_\lambda) \in \omega(\log \lambda)$. Suppose that Circ is two-point pseudorandomness-style puncturable for $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}^{\mathsf{pr}}_{\mathsf{ID}_{\mathcal{D}}})$ and that Π satisfies $\mathsf{ID}_{\mathcal{D}}$-generalized UPO security and $i\mathcal{O}$-security. Then CP_Π satisfies oracular pseudorandomness-style copy-protection security for $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}^{\mathsf{pr}}_{\mathsf{ID}_{\mathcal{D}}})$.*

Proof. Specialize the hybrids of [ABH+26, Theorem 7.4] to $x_{\mathcal{B}} = x_{\mathcal{C}} = x \leftarrow \mathcal{D}_\lambda$ and $b_{\mathcal{B}} = b_{\mathcal{C}} = b$. The programming and puncturing hybrids use one shared uniform value. Since the protected circuit and C differ only at x, the punctured oracles given to $\mathcal{B}$ and $\mathcal{C}$ are identical in every corresponding hybrid, and hence the cross-point oracle switches of [ABH+26, Claim 7.5] are unnecessary.

It remains to justify the oracle switch for $\mathcal{A}$. Let $q = q(\lambda)$ be a polynomial upper bound on the number of its oracle queries and let W be their total query weight on x immediately before this switch. The case $q = 0$ is immediate, so suppose $q \geq 1$. In the initial experiment, x is sampled

independently of $\mathcal{A}$'s pre-challenge view. Thus, if an index is sampled uniformly from $[q]$ and the input of that query is measured to obtain x', with $x' := \bot$ if the query does not occur, then

$$\Pr\left[x' = x\right] \leq 2^{-\widetilde{H}_\infty(\mathcal{D}_\lambda)} = \mathsf{negl}(\lambda). \tag{132}$$

The preceding programming and puncturing hybrids remain indistinguishable when this measurement is included. In the hybrid immediately before the oracle switch, the same test succeeds with probability $\mathbb{E}[W]/q$. Therefore, exactly as in [ABH+26, Claim 7.6],

$$\mathbb{E}[W] \leq q \cdot 2^{-\widetilde{H}_\infty(\mathcal{D}_\lambda)} + \mathsf{negl}(\lambda) = \mathsf{negl}(\lambda). \tag{133}$$

The BBBV query bound [ABH+26, Theorem 2.1] then changes the oracle between C and the programmed circuit at negligible cost. The last $i\mathcal{O}$ hybrid replaces the programmed circuit by the functionally equivalent circuit produced by the generalized-puncturing algorithm. The resulting experiment is exactly the $\mathsf{ID}_{\mathcal{D}}$-generalized UPO experiment, and its winning probability is at most $1/2 + \mathsf{negl}(\lambda)$. □

We remark that for independent high-min-entropy challenge points with a common challenge bit, the corresponding oracular pseudorandomness-style result is already shown in [ABH+26, Theorem 7.4], and the oracular unpredictability-style result follows from [ABH+26, Theorem 8.1 and Remark 8.3].

Next Theorem 30 combined with our feasibility result for UPO (Theorem 23) gives us the following immediate corollaries.

Corollary 10 (Oracular pseudorandomness-style copy protection for identical challenges)**.** *Assume post-quantum polynomially secure $i\mathcal{O}$ and the quantum hardness of LWE. Let $c > 0$ be a constant, let $\mathcal{D}$ be QPT-samplable with $\widetilde{H}_\infty(\mathcal{D}_\lambda) \geq \lambda^c$, and suppose that Circ is two-point pseudorandomness-style puncturable for $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}^{\mathsf{pr}}_{\mathsf{ID}_{\mathcal{D}}})$. Then there exists a copy-protection scheme for Circ satisfying oracular pseudorandomness-style security for $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}^{\mathsf{pr}}_{\mathsf{ID}_{\mathcal{D}}})$.*

Proof. Combine Theorems 18 and 23 and Lemma 6 with Theorem 30. □

For the OWF-based instantiation, we again retain identical-uniform challenges.

Corollary 11 (Identical-uniform oracular pseudorandomness-style copy protection)**.** *Assume post-quantum polynomially secure $i\mathcal{O}$ and post-quantum one-way functions. Let $\ell_{\mathsf{in}} \in \omega(\log \lambda)$, and suppose that Circ is two-point pseudorandomness-style puncturable for $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}^{\mathsf{pr}}_{\mathsf{ID}_{\mathsf{U}}})$. Then there exists a copy-protection scheme for Circ satisfying oracular pseudorandomness-style security for $(\mathcal{D}_{\mathsf{Circ}}, \mathsf{Samp}^{\mathsf{pr}}_{\mathsf{ID}_{\mathsf{U}}})$.*

Proof. The "Moreover" part of Theorem 23 gives ID_{U}-generalized UPO security. Add $i\mathcal{O}$-security using [ABH+26, Lemma 3.16], and apply Theorem 30 with $\mathcal{D}_\lambda$ uniform over $\mathcal{X}_\lambda$. □

Corollary 12 (Oracular pseudorandomness-style copy protection for PRFs)**.** *Assume post-quantum polynomially secure $i\mathcal{O}$ and post-quantum one-way functions. Then there exists a puncturable PRF family on every domain $\{0,1\}^{\ell_{\mathsf{in}}(\lambda)}$ with $\ell_{\mathsf{in}} \in \omega(\log \lambda)$ having a copy-protection scheme satisfying oracular pseudorandomness-style security for $\mathsf{Samp}^{\mathsf{pr}}_{\mathsf{ID}_{\mathsf{U}}}$.*

Proof. Post-quantum one-way functions imply post-quantum puncturable PRFs. Since the identical sampler yields the singleton set $\mathcal{S} = \{x\}$, ordinary puncturable-PRF security gives Definition 43; apply Corollary 11. □

References

[Aar09] Scott Aaronson. Quantum copy-protection and quantum money. In *24th Annual IEEE Conference on Computational Complexity*, pages 229–242, 2009.

[Aar16] Scott Aaronson. The complexity of quantum states and transformations: From quantum money to black holes, 2016.

[AB24] Prabhanjan Ananth and Amit Behera. A modular approach to unclonable cryptography. In *Advances in Cryptology – CRYPTO 2024*, page 3–37. Springer Nature Switzerland, 2024.

[ABH+26] Prabhanjan Ananth, Amit Behera, Zikuan Huang, Fuyuki Kitagawa, and Takashi Yamakawa. Copy-protection from unclonable puncturable obfuscation, revisited. In *Advances in Cryptology – EUROCRYPT 2026*, page 541–570. Springer Nature Switzerland, 2026.

[AC12] Scott Aaronson and Paul Christiano. Quantum money from hidden subspaces. Cryptology ePrint Archive, Report 2012/171, 2012.

[AK21] Prabhanjan Ananth and Fatih Kaleoglu. Unclonable encryption, revisited. In Kobbi Nissim and Brent Waters, editors, *TCC 2021: 19th Theory of Cryptography Conference, Part I*, volume 13042 of *Lecture Notes in Computer Science*, pages 299–329, Raleigh, NC, USA, November 8–11, 2021. Springer, Cham, Switzerland.

[AKL+22] Prabhanjan Ananth, Fatih Kaleoglu, Xingjian Li, Qipeng Liu, and Mark Zhandry. On the feasibility of unclonable encryption, and more. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part II*, volume 13508 of *Lecture Notes in Computer Science*, pages 212–241, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland.

[AKL23] Prabhanjan Ananth, Fatih Kaleoglu, and Qipeng Liu. Cloning games: A general framework for unclonable primitives. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part V*, volume 14085 of *Lecture Notes in Computer Science*, pages 66–98, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland.

[AKY25a] Prabhanjan Ananth, Fatih Kaleoglu, and Henry Yuen. Simultaneous haar indistinguishability with applications to unclonable cryptography. In Raghu Meka, editor, *ITCS 2025: 16th Innovations in Theoretical Computer Science Conference*, volume 325, pages 7:1–7:23, New York, NY, USA, January 7–10, 2025. Leibniz International Proceedings in Informatics (LIPIcs).

[AKY25b] Prabhanjan Ananth, Fatih Kaleoglu, and Henry Yuen. Simultaneous haar indistinguishability with applications to unclonable cryptography. In Raghu Meka, editor, *16th Innovations in Theoretical Computer Science Conference, ITCS 2025, Columbia University, New York, NY, USA, January 7-10, 2025*, volume 325 of *LIPIcs*, pages 7:1–7:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025.

[ALL+21] Scott Aaronson, Jiahui Liu, Qipeng Liu, Mark Zhandry, and Ruizhe Zhang. New approaches for quantum copy-protection. In Tal Malkin and Chris Peikert, editors,

Advances in Cryptology – CRYPTO 2021, Part I, volume 12825 of *Lecture Notes in Computer Science*, pages 526–555, Virtual Event, August 16–20, 2021. Springer, Cham, Switzerland.

[AS26] Prabhanjan Ananth and Amit Sahai. Unconditional unclonable encryption, 2026.

[BBC26] Archishna Bhattacharyya, Anne Broadbent, and Eric Culf. Statistically secure uncloneable encryption of arbitrary messages, 2026.

[BDF26] Pierre Botteron, Sébastien Designolle, and Omar Fawzi. Pauli encodings & unclonable encryption, 2026.

[BL20] Anne Broadbent and Sébastien Lord. Uncloneable quantum encryption via oracles. volume 158, pages 4:1–4:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020.

[BPW16] Nir Bitansky, Omer Paneth, and Daniel Wichs. Perfect structure on the edge of chaos - trapdoor permutations from indistinguishability obfuscation. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A: 13th Theory of Cryptography Conference, Part I*, volume 9562 of *Lecture Notes in Computer Science*, pages 474–502, Tel Aviv, Israel, January 10–13, 2016. Springer Berlin Heidelberg, Germany.

[ÇG26] Alper Çakan and Vipul Goyal. How to copy-protect malleable-puncturable cryptographic functionalities under arbitrary challenge distributions: A unified solution to quantum protection. In *Advances in Cryptology – EUROCRYPT 2026*, pages 481–509. Springer Nature Switzerland, 2026.

[CGLZR23] Alper Cakan, Vipul Goyal, Chen-Da Liu-Zhang, and João Ribeiro. Unbounded leakage-resilience and intrusion-detection in a quantum world. Cryptology ePrint Archive, Report 2023/410, 2023.

[CHV23] Céline Chevalier, Paul Hermouet, and Quoc-Huy Vu. Towards unclonable cryptography in the plain model. Cryptology ePrint Archive, Report 2023/1825, 2023.

[CLLZ21] Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. Hidden cosets and applications to unclonable cryptography. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 556–584, Virtual Event, August 16–20, 2021. Springer, Cham, Switzerland.

[CMP20] Andrea Coladangelo, Christian Majenz, and Alexander Poremba. Quantum copy-protection of compute-and-compare programs in the quantum random oracle model. Cryptology ePrint Archive, Report 2020/1194, 2020.

[CMP24] Andrea Coladangelo, Christian Majenz, and Alexander Poremba. Quantum copy-protection of compute-and-compare programs in the quantum random oracle model. *Quantum*, 8:1330, May 2024.

[CV22] Eric Culf and Thomas Vidick. A monogamy-of-entanglement game for subspace coset states. *Quantum*, 6:791, 2022.

[GM82] Shafi Goldwasser and Silvio Micali. Probabilistic encryption and how to play mental poker keeping secret all partial information. In *14th Annual ACM Symposium on*

Theory of Computing, pages 365–377, San Francisco, CA, USA, May 5–7, 1982. ACM Press.

[GM84] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.

[Gol01] Oded Goldreich. *Foundations of Cryptography*. Cambridge University Press, 2001.

[Gol04] Oded Goldreich. *Foundations of Cryptography*. Cambridge University Press, 2004.

[GZ20] Marios Georgiou and Mark Zhandry. Unclonable decryption keys. Cryptology ePrint Archive, Report 2020/877, 2020.

[KN23] Fuyuki Kitagawa and Ryo Nishimaki. One-out-of-many unclonable cryptography: Definitions, constructions, and more. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023: 21st Theory of Cryptography Conference, Part IV*, volume 14372 of *Lecture Notes in Computer Science*, pages 246–275, Taipei, Taiwan, November 29 – December 2, 2023. Springer, Cham, Switzerland.

[KY25] Fuyuki Kitagawa and Takashi Yamakawa. Foundations of single-decryptor encryption. In *Theory of Cryptography Conference*, pages 166–197. Springer, 2025.

[NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.

[SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th Annual ACM Symposium on Theory of Computing*, pages 475–484, New York, NY, USA, May 31 – June 3, 2014. ACM Press.

[TFKW13] Marco Tomamichel, Serge Fehr, Jedrzej Kaniewski, and Stephanie Wehner. One-sided device-independent QKD and position-based cryptography from monogamy games. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, volume 7881 of *Lecture Notes in Computer Science*, pages 609–625, Athens, Greece, May 26–30, 2013. Springer Berlin Heidelberg, Germany.

[Tro12] Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.

[Wat06] John Watrous. Zero-knowledge against quantum attacks. In Jon M. Kleinberg, editor, *38th Annual ACM Symposium on Theory of Computing*, pages 296–305, Seattle, WA, USA, May 21–23, 2006. ACM Press.

[Wie83] Stephen Wiesner. Conjugate coding. *ACM SIGACT News*, 15(1):78–88, 1983.

[Win99] Andreas Winter. Coding theorem and strong converse for quantum channels. *IEEE Transactions on Information Theory*, 45(7):2481–2485, 1999.

[Zha12] Mark Zhandry. How to construct quantum random functions. In *53rd Annual Symposium on Foundations of Computer Science*, pages 679–687, New Brunswick, NJ, USA, October 20–23, 2012. IEEE Computer Society Press.

[Zha20] Mark Zhandry. Schrödinger's pirate: How to trace a quantum decoder. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020: 18th Theory of Cryptography Conference, Part III*, volume 12552 of *Lecture Notes in Computer Science*, pages 61–91, Durham, NC, USA, November 16–19, 2020. Springer, Cham, Switzerland.

[Zha21] Mark Zhandry. Quantum lightning never strikes the same state twice. or: Quantum money from cryptographic assumptions. *J. Cryptol.*, 34(1):6, 2021.