



Consistency for Functional Encryption

Christian Badertscher¹ , Aggelos Kiayias^{1,2},
Markulf Kohlweiss^{1,2}, and Hendrik Waldner² 

¹ IOHK

{christian.badertscher,aggelos.kiayias}@iohk.io

² University of Edinburgh, Edinburgh, UK

{markulf.kohlweiss,hendrik.waldner}@ed.ac.uk

Abstract. In functional encryption (FE) a sender, Alice, encrypts plaintexts that a receiver, Bob, can obtain functional evaluations of, while Charlie is responsible for initializing the encryption keys and issuing the decryption keys. Standard notions of security for FE deal with a malicious Bob and how the confidentiality of Alice’s messages can be maintained taking into account the leakage that occurs due to the functional keys that are revealed to the adversary via various forms of indistinguishability experiments that correspond to IND-CPA, IND-CCA and simulation-based security.

In this work we provide a complete and systematic investigation of *Consistency*, a natural security property for FE, that deals with attacks that can be mounted by Alice, Charlie or a collusion of the two against Bob. We develop three main types of consistency notions according to which set of parties is corrupted and investigate their relation to the standard security properties of FE. To validate our different consistency types, we investigate FE in the universally composition setting and we show that our consistency notions naturally complement FE security by proving how they imply (and are implied by) UC security depending on which set of parties is corrupted; in this way we demonstrate a complete characterization of consistency for FE. Finally, we provide explicit constructions that achieve consistency efficiently either directly via a construction based on MDDH for specific function classes of inner products over a modulo group or generically for all the consistency types via compilers using standard cryptographic tools.

1	Introduction	1
1.1	Contributions of this Work	2
1.2	Comparison with Related Work	3
2	Preliminaries	4
2.1	Functional Encryption	5
2.2	Security Definitions	6
2.3	Standard Tools and Assumptions	7
2.4	Inner-product Functionality Classes	8
2.5	Non-interactive Proofs	8
2.6	Verifiable Functional Encryption	10
2.7	Universal Composability	10
3	Consistency for Functional Encryption Schemes	11
3.1	Consistency with a dishonest Input Provider	11
3.2	Consistency with a dishonest Input Provider and Key Generator	14
3.3	Consistency with a dishonest Parameter/Key Generator	15
4	Relations (in)between Consistency and Confidentiality	17
4.1	Relations among the Consistency Notions	17
4.2	Consistency does not imply Confidentiality	18
4.3	Confidentiality does not imply Consistency	19
4.4	Consistency does not amplify Confidentiality	22
5	Consistency Analysis of Selected Functional Encryption Schemes	23
5.1	Inconsistency of the Plain Schemes	23

5.2	Consistency for Inner-product Schemes	24
5.3	Consistency of a related Predicate Scheme	26
6	Consistency Compilers	28
6.1	Input Consistency	28
6.2	Setup Consistency Compilers	37
6.3	Strong Input Consistency	46
7	UC Consistency for Functional Encryption	48
	Acknowledgments	53
A	Overview of the UC Framework	57
B	Consistency for Different Types of Functional Encryption	58
C	Details Of the UC Analysis	58
C.1	Assumed Functionalities	58
C.2	Proof of the UC Realization (Theorem 7.2)	59

1 Introduction

Functional encryption (FE) [23, 56] has emerged as an important and general purpose cryptographic primitive, extending and generalizing earlier more specialized encryption concepts that include Identity-Based Encryption [22], Attribute-Based Encryption [45, 62] and Predicate Encryption [50]. Similar to these earlier primitives, in FE, there exists a setup algorithm that produces a master public-key mpk and a master secret-key msk , and a key-generation algorithm that receives as input msk and a function f and produces a function-specific secret-key sk_f . Subsequently, using sk_f along with the decryption algorithm, the computation of the value $f(x)$ is facilitated given any ciphertext that encrypts x . The potential applications of FE are numerous and include any setting where there exist designated entities that are entitled to functional views of encrypted information that is described in the form of a function f for which an associated functional key sk_f is produced by the key-generation procedure.

In order to define correctness and security of FE it is helpful to identify three distinct entities associated with the algorithms that comprise any FE scheme. Alice is the sender, wishing to transmit data x , Bob is a recipient wishing to receive $f(x)$ for some function $f(\cdot)$ and Charlie is an authority that issues the (master) keys. Typically we think there are multiple Alice and Bob parties for any given setup instance created by Charlie. As one of these Bob parties can be corrupted, this also captures security against an eavesdropper that only observes the network. Correctness mandates the natural requirement that Bob receives the value $f(x)$ for properly encrypted ciphertexts prepared by Alice that contain x . Security on the other hand is typically captured as a game with an adversary who attempts to distinguish between two possible plaintexts x_0, x_1 for which it holds that $f_i(x_0) = f_i(x_1)$ for all functions f_i whose key is possessed by the adversary. A stronger notion of security puts forth a simulation-based formulation and asks that ciphertexts can be simulated in an indistinguishable way. Cf. [6, 7, 15, 17, 23, 33, 44, 48, 53, 56]. The adversary controls multiple different Bob sessions and typically interferes with the honest Alice only in the sense of chosen plaintext attacks, however chosen ciphertext attacks have also been considered [19] (in which case the adversary may e.g., manipulate Alice’s ciphertext and submit it to Bob’s decryption oracle). Specifically, in this work we build on the composable formalization of Matt and Maurer [53].

Consistency problems in real-world applications of FE. A crucial problem for any cryptographic primitive is to identify the exact set of correctness and security properties that are necessary and sufficient for deploying the primitive within an intended real world system. To see that there is a fundamental property of FE that is missing, it is helpful to recall the most well known applications of FE and showcase the problems that emerge when *consistent* behavior of an FE scheme is not guaranteed.

Processing Encrypted Data. In [23] the following motivating example for FE is presented: Alice encrypts a photograph x and uploads to her cloud service provider while Bob, a law enforcement agent, wishes to check whether any photographs in the cloud match a specific face. Using FE, Bob can achieve his objective, taking advantage of a functional key which detects the encrypted photographs which match the face being searched without revealing any information about anything else. Given the above setting, it is in everyone’s understanding that if a photograph contains someone that matches the specific face being searched, the law enforcement agent will be able to detect it. Nevertheless, neither standard notions of security nor correctness of FE can rule out the possibility that a malicious Alice creates a ciphertext that will be misclassified by Bob, specifically, a ciphertext that decrypts to a photograph of the person being searched, and for which the employed face recognition algorithm f works, but is not detected as such by Bob when employing sk_f . The same “misclassification” inconsistency issue applies to any setting where FE is used to classify ciphertexts in-transit or in-situ (e.g., for virus-detection, routing etc.).

Attribute Based Encryption. In an attribute-based encryption (ABE) scheme [45], which is a special case of FE, Alice encrypts a message together with a set of attributes γ . Subsequently, Bob, who possesses a key corresponding to an access structure \mathbb{A} will be able to decrypt the message as long as $\gamma \in \mathbb{A}$. Consider now also another party, say Bob junior, possessing a key for the access structure $\mathbb{A}' \subsetneq \mathbb{A}$. Given the above setting, it is in everyone’s understanding that whatever messages Bob junior is able to see, Bob should see as well.

Nevertheless, neither standard notions of security nor correctness of FE can rule out the possibility that a malicious Alice crafts a ciphertext that Bob junior will be able to decrypt but Bob would not.

Consistency as a fundamental property for FE. What do the above problems tell us? Similar to advanced properties of ordinary PKE (such as e.g., robustness [1]), advanced properties for FE are needed when using the primitive in a real world setting because such properties are implied by the way the primitive is understood in the real-world. Moreover, the level at which they should be defined is at the level of the basic definition and syntax of FE. We call the enhanced property the above issues point to *consistency*; it addresses, at minimum, the adversarial setting where a malicious Alice produces a specially crafted ciphertext that causes an honest Bob to misclassify it, or, perhaps even a malicious Charlie who tampers with the setup to cause further types of misclassification. Interestingly and somewhat surprisingly, such a consistency property has not been considered in the strict context of FE so far and enhanced FE schemes, departing from the standard syntax such as [12], do only consider certain consistency aspects (see below). We show that, as with the confidentiality of FE, the consistency of FE has several flavors, some of which are very efficient to ensure, while others require more sophisticated techniques.

1.1 Contributions of this Work

We roll out consistency as a fundamental property for a generic FE scheme from first principles. We provide a number of constructions for various consistency and security notions either directly for specific function classes or generically via compilers that upgrade existing FE constructions to be consistent. To formally cross-check our new notion, we show that the defined properties are necessary and sufficient in realizing the UC characterization of an “ideally” secure and consistent FE-scheme abstraction derived from [53]. The modelling of all security properties as an ideal functionality assures that no important details were omitted and that our game-based definitions interoperate correctly. In more details we make the following contributions.

Formal definition of consistency. We identify three main types of consistency, each type naturally corresponding to a particular set of corrupted parties. The formalization is given in Section 3.

- *Input consistency* considers a malicious Alice who computes a ciphertext ct as well as some candidate functions f_i . The ciphertext ct is decrypted under sk_{f_i} to obtain the values y_i and the adversary wins if there is no single x that can explain ct in the sense that $f_i(x) = y_i$. We also incorporate in the definition of input consistency a concept of extractability that facilitates the relation of the primitive to the universal composable security of FE (see below).
- *Strong input consistency* couples the above goal with additional adversarial power, by considering the setting where both Alice and Charlie are corrupted, and therefore, subverted parameters can assist the adversary in breaking the scheme.
- *Setup consistency* is the consistency notion that deals with a malicious Charlie. In this setting the adversary issues two plaintexts x_1, x_2 as well as a secret-key and a function f . The plaintexts are honestly encrypted and subsequently their decryptions y_1, y_2 are evaluated. The adversary wins the game if exactly one of the decryptions fails or $y_i \neq f(x_i)$ for some i . While at first sight it seems that setup consistency is implied by e.g. strong input consistency, we point out that this is not the case.

We highlight that consistency in the above sense complements security, as in the latter Alice and Charlie are honest and Bob is malicious. To show that our definitions do formally capture what they are intended for, we put forth in Section 7 a complete treatment of consistent FE in the universal composition (UC) setting [26]. Specifically, we prove that input consistency/setup consistency/strong input consistency is sufficient and necessary for UC security in the case Alice/Charlie/Alice+Charlie are corrupted respectively. This pairs and complements the result of [53] which implies that CFE security is sufficient (and necessary) for UC security in the case Bob is corrupted. This completes the positioning of consistency as an important novel property of FE.

Systematic study of consistency vis-à-vis existing security properties. We carefully analyze the relations in-between the consistency notions and between consistency and security. We confirm our intuition that all notions define separate levels of consistency, the only exception being that strong input consistency implies input consistency. With respect to security, namely IND-CPA, IND-CCA and CFE, which is composable security notion for FE, we show that strong input consistency does not imply IND-CPA security constructively. Furthermore, we show that IND-CPA together with strong input consistency does not imply any of the other stronger security notions such as IND-CCA or CFE. Finally, IND-CCA and CFE individually do not imply input consistency. We refer to Figure 8 for a relation diagram. Thus, it follows that consistency is independent from existing notions of FE security. The proofs are given in Section 4.

Realizing FE with consistency. We first describe, in Section 5, concrete input-consistent constructions for an inner-product type of FE under the Matrix DDH assumption for two different functional classes. The first construction covers the modified inner-product functionality class over a modulo group and the second construction covers the function class of exponentiated inner-products over a modulo group. Both of these constructions are adapted from the construction of [8]. Interestingly, we prove that previous efficient constructions for the function class of inner-product over the *integers* all fail to provide input consistency (we present explicit attacks).

Subsequently, in Section 6 we present several compilers that achieve consistency in a black-box manner from an FE scheme:

- For input consistency we use (designated verifier) NIZKs to compile any FE scheme into a corresponding one that satisfies input consistency under standard assumptions. We show that the construction preserves CPA and CFE security and provide an additional construction to *lift* the privacy to obtain a CCA-secure input-consistent FE scheme only assuming a CPA secure FE scheme under the same assumptions.
- For setup consistency we use a form of the twin encryption technique [55] and NIWIs, [14, 21, 47], which allows Charlie to demonstrate that the secret-keys are properly generated. Here again, we provide constructions that preserve CPA security (and CFE under certain assumptions) and provide an advanced compiler that lifts an assumed CPA-secure FE scheme to a setup-consistent CCA-secure FE scheme.
- Finally, for strong input consistency, we establish the formal connection to *verifiable FE* schemes as introduced in [12] and show that any such scheme can be turned into a strong input consistent FE scheme and thus one can deploy the compiler of [12] to extract a scheme from standard FE. The reverse, however is not true and we elaborate on this in Section 6 as well as below in Section 1.2.

1.2 Comparison with Related Work

Relation to robustness notions. As a first approximation, input-consistency for FE can be thought of as a natural well-formedness property of FE ciphertexts which is the main reason why it is of relevance to the cryptographic investigation of the FE primitive, arguably in the same way that other types of consistency properties for regular public-key encryption are. For instance, plaintext-awareness [32] and non-malleability [34] are security properties for public-key encryption that deal with ciphertext well-formedness and are independent of weaker notions of security such as indistinguishability against chosen plaintext attacks while related to stronger formulations of security such as indistinguishability against chosen ciphertext attacks. Furthermore, a consistency-like property more closely related to FE is robustness of identity-based encryption (IBE) [1, 36]. In the strong robustness attack of Abdalla et al. [1] against IBE an adversary outputs two identities $\text{id}_1 \neq \text{id}_2$ and a ciphertext ct . The game derives decryption keys for both identities and the adversary wins if the decryption of ct is non- \perp under both keys. IBE can be viewed as a special case of FE: encrypt the pair identity and message and let the user with identity id possess the key for the function “ $f_{\text{id}}(\text{id}', m) = m$, if $\text{id} = \text{id}'$ else \perp .” It is immediate that a robustness attack is a consistency attack against the above FE scheme, since it cannot be that two distinct identities id_1, id_2 equal the same id' . By this reduction, we see that our notion of consistency for general FE can be instantiated—notably not only for IBE—to yield such related robustness notions for special cases directly. Besides strong robustness, the authors also introduce the notion of weak robustness, in which the adversary outputs a message and two identities and the challenger encrypts

the message under the first identity and tries to decrypt it using the second identity. The authors show that weak robustness is implied by strong robustness, and since strong robustness mirrors our notion of input consistency, it follows immediately that weak robustness is also implied by input consistency notion when instantiated for IBE as above.

A related kind of work is the task of generalizing the traditional notions of robustness, which roughly captures that decryption with a secret key that is generated in some system A *must indicate a failure* when presented with a ciphertext that was generated using (different) parameters of some system B . These notions have been extended (also to FE) [36, 40], but none of them looks at the much harder problem, namely to ensure that decrypted values (for example within one system), make sense relatively to each other, especially in FE, which we tackle in this work.

Relation to verifiable FE. Prior to our work, there was only one previous and very insightful work [12] which identified some of the above deficiencies and, to address them, put forth a cryptographic primitive which is substantially stronger than FE, called verifiable functional encryption (VFE) (and the compiler presented in [12] has recently been instantiated using pairing-based NIWIs and a perfectly correct functional encryption scheme for predicates over inner-products [63]). VFE extends the normal FE syntax by two additional predicates to check validity of keys and ciphertexts, respectively. As already mentioned above, VFE implies strong input-consistent FE but interestingly the reverse is not necessarily true as VFE requires the public verifiability of ciphertext and functional key well-formedness whereas for strong input consistency a private-key based test, for instance, would suffice.

Furthermore, VFE (as well as strong input-consistent FE) does not imply setup consistency, since it merely guarantees that an encryption c of plaintext x would consistently decrypt to something but *not necessarily* to functions of x that an honest sender has encrypted (i.e., the setting where genuinely generated ciphertexts may be mangled due to a subverted setup). This, however, seems rather crucial, as additional guarantees for this setting where only Charlie is dishonest might be desirable in various settings (see Section 3). Finally, we point out that aiming at input or setup consistency allows for more efficient constructions, compared to the verifiable FE compiler presented in [12]. In the case of input consistency, the compiler we present only relies on NIZKs (not NIWIs) and only requires a single instance of a functional encryption scheme (instead of four). For the setup compiler, we also require NIWIs, but only three instances of the functional encryption scheme instead of four.

Relation to distributed setup-generation. One can view setup consistency as focusing on the important question of setup subversion resistance (again w.r.t. misclassification attacks for honestly generated ciphertexts in FE). A very different approach, namely generating setup parameters or keys in a decentralized fashion as for example in [51], requires changes in the syntax of FE as multiple parties need to participate. But even in this setting, our setup consistency notion additionally gives guarantees if all the parties (or a number of parties exceeding a certain critical threshold) involved in the MPC protocol are corrupted. The two approaches to securing the setup and countering subversion are thus orthogonal by nature.

2 Preliminaries

General Notation. We denote, the security parameter with $\lambda \in \mathbb{N}$ and use 1^λ as its unary representation. We call a randomized algorithm \mathcal{A} *probabilistic polynomial time* (PPT), if there exists a polynomial $p(\cdot)$ such that for every input x the running time of $\mathcal{A}(x)$ is bounded by $p(|x|)$. A function $\text{negl} : \mathbb{N} \rightarrow \mathbb{R}^+$ is called *negligible* if for every positive polynomial $p(\lambda)$ a $\lambda_0 \in \mathbb{N}$ exists, such that for all $\lambda > \lambda_0 : \epsilon(\lambda) < 1/p(\lambda)$. The set $\{1, \dots, n\}$ is denoted as $[n]$ for $n \in \mathbb{N}$. For the equality check of two elements, we use “=”. The assign operator is denoted with “:=”, whereas randomized assignment is denoted with $a \leftarrow A$, with a randomized algorithm A and where the randomness is not explicit. If the randomness is explicit we write $a := A(x; r)$ where x is the input and r is the randomness. For algorithms \mathcal{A} and \mathcal{B} , we write $\mathcal{A}^{\mathcal{B}(\cdot)}(x)$ to denote that \mathcal{A} gets x as an input and has oracle access to \mathcal{B} , that is, the response for an oracle query q is $\mathcal{B}(q)$. We use $\mathcal{A}(\cdot)[[s]]$ to denote that \mathcal{A} gets an additional input s which it can update. In more detail, $\mathcal{A}(x)[[s]]$ corresponds to the algorithm that invokes $(y, s) \leftarrow \mathcal{A}(x, s)$ and returns y .

We write e_i^ℓ for the unit vector of length ℓ that is 1 at position i and 0 everywhere else. We omit the length when it is clear from the context.

For the generation of prime-order groups, let GGen be a PPT algorithm that on input 1^λ returns a description $\mathcal{G} = (\mathbb{G}, p, g)$ of a cyclic group \mathbb{G} of order p for a λ -bit prime p , whose generator is g . We use the implicit representation $[x]_g$ for group elements of the form g^x with a generator $g \in \mathbb{G}$. This notation is also used in the case of matrices. In more detail, for a matrix $\mathbf{A} = (a_{i,j}) \in \mathbb{Z}_p^{n \times m}$, we define $[\mathbf{A}]_g$ as the implicit representation of \mathbf{A} in \mathbb{G} :

$$[\mathbf{A}]_g := \begin{pmatrix} g^{a_{1,1}} & \dots & g^{a_{1,m}} \\ \vdots & & \vdots \\ g^{a_{n,1}} & \dots & g^{a_{n,m}} \end{pmatrix}$$

2.1 Functional Encryption

We now introduce the relevant notation for functional encryption.

Definition 2.1. We denote by $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ a family of sets \mathcal{F}_λ of functions $f : \mathcal{X}_\lambda \rightarrow \mathcal{Y}_\lambda$. We call \mathcal{F}_λ a *functionality class* such that all functions $f \in \mathcal{F}_\lambda$ have the same domain and the same range. We omit λ when it is clear from the context.

For notational convenience, we further define an extension for functions $f \in \mathcal{F}$ in order to develop a formal language that simplifies expressing consistency failures later in this work. We introduce two additional error symbols \perp, \diamond and formally include them in the domain or range of the functions as defined below. We note that both symbols do not have any influence on the behavior of the function f . Rather, we require that the symbol \perp maps to \perp and that symbol \diamond has no preimage:

Definition 2.2 (Function Extension). Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a function of the functionality class \mathcal{F} , we define a function $\tilde{f} : (\mathcal{X} \cup \{\perp\}) \rightarrow (\mathcal{Y} \cup \{\perp, \diamond\})$, with $\perp, \diamond \notin \mathcal{X}, \mathcal{Y}$. The function \tilde{f} has the following behavior:

$$\tilde{f}(x) = \begin{cases} f(x) & \text{if } x \in \mathcal{X} \\ \perp & \text{if } x = \perp \end{cases} \quad \text{and} \quad \tilde{f}^{-1}(y) = \begin{cases} f^{-1}(y) & \text{if } y \in \mathcal{Y} \\ \{\perp\} & \text{if } y = \perp \\ \emptyset & \text{if } y = \diamond \end{cases}.$$

For a (standard) functionality class \mathcal{F} , the induced extended class is the set of function extensions of all $f \in \mathcal{F}$. When clear from the context, we do not introduce a new symbol for the extended class.

A functional encryption scheme is defined in the following way, where we follow the syntax of [12].

Definition 2.3 (Functional Encryption). Let $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ be a family of sets \mathcal{F}_λ of functions $f : \mathcal{X}_\lambda \rightarrow \mathcal{Y}_\lambda$ such that \mathcal{F}_λ contains a distinguished leakage function³ f_0 . A functional encryption scheme (FE) for the functionality class \mathcal{F}_λ is a tuple of four algorithms $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$:

Setup(1^λ): Takes as input a unary representation of the security parameter λ and outputs the master public key mpk and the master secret key msk .

KeyGen($\text{mpk}, \text{msk}, f$): Takes as input the master public mpk , the master secret key msk and a function $f \in \mathcal{F}_\lambda$, and outputs a functional key sk_f . The key for the leakage function f_0 is the empty string denoted by ε .

Enc(mpk, x): Takes as input the master public key mpk and a string $x \in \mathcal{X}_\lambda$, and outputs a ciphertext ct or err (to denote an encryption error).

Dec($\text{mpk}, f, \text{sk}_f, \text{ct}$): Takes as input a functional key sk_f and a ciphertext ct and outputs a function value $y \in \mathcal{Y}_\lambda$ or one of the special symbols of the function extension: \perp indicates an invalid ciphertext and \diamond invalid keys.

³ The leakage function is a modelling technique adopted from [23] that can, e.g., reveal information about the plaintext length.

A scheme FE is correct, if (for all $\lambda \in \mathbb{N}$), for all pairs (mpk, msk) in the support of $\text{Setup}(1^\lambda)$ all functions $f \in \mathcal{F}_\lambda$ and input values $x \in \mathcal{X}_\lambda$, it holds that

$$\Pr[\text{Dec}(\text{mpk}, f, \text{KeyGen}(\text{mpk}, \text{msk}, f), \text{Enc}(\text{mpk}, x)) = f(x)] = 1.$$

For notational simplicity, we omit certain input values when they are not required by a concrete scheme (such as the additional mpk or f when decrypting).

The security of functional encryption is formally captured by the CPA [23], CCA2 [19], as well as the CFE [53] (composable) security notions, that formalize, roughly speaking that an attacker does not learn anything beyond what he can anyway decrypt given a set of functional keys he requested (and in the case of CCA, even if an additional decryption oracle is available).

2.2 Security Definitions

Definition 2.4 (CPA & CCA Security of FE). Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be a functional encryption scheme, $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ a function family and $\beta \in \{0, 1\}$. We define the experiments $\text{IND-CPA}_\beta^{\text{FE}}(1^\lambda, \mathcal{A})$ and $\text{IND-CCA}_\beta^{\text{FE}}(1^\lambda, \mathcal{A})$ in Fig. 1. The associated advantage of an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ for $\text{XX} \in \{\text{CPA}, \text{CCA}\}$ is defined by

$$\text{Adv}_{\text{FE}, \mathcal{A}}^{\text{IND-XX}}(\lambda) = |\Pr[\text{IND-XX}_0^{\text{FE}}(1^\lambda, \mathcal{A}) = 1] - \Pr[\text{IND-XX}_1^{\text{FE}}(1^\lambda, \mathcal{A}) = 1]|.$$

An adversary \mathcal{A} is valid if for the two submitted challenges x^0 and x^1 and all keys sk_f the attacker obtained for f via calls to KeyGen (and including the empty key for f_0), it holds that $f(x^0) = f(x^1)$. For CCA security, the adversary \mathcal{A} is additionally not allowed to query the decryption oracle $\text{QDec}(f, \text{ct})$ on the obtained challenge ciphertext $\text{ct} = \text{Enc}(\text{mpk}, x^\beta)$.

A functional encryption scheme FE is IND-XX secure, if for any valid PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, there exists a negligible function negl , such that $\text{Adv}_{\text{FE}, \mathcal{A}}^{\text{IND-XX}}(\lambda) \leq \text{negl}(\lambda)$.

$\text{IND-CPA}_\beta^{\text{FE}}(1^\lambda, \mathcal{A})$	$\text{IND-CCA}_\beta^{\text{FE}}(1^\lambda, \mathcal{A})$
$(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$	$(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$
$(x^0, x^1, \text{st}) \leftarrow \mathcal{A}_1^{\text{KeyGen}(\text{mpk}, \text{msk}, \cdot)}(\text{mpk})$	$(x^0, x^1, \text{st}) \leftarrow \mathcal{A}_1^{\text{KeyGen}(\text{mpk}, \text{msk}, \cdot), \text{QDec}(\cdot, \cdot)}(\text{mpk})$
$\text{ct} \leftarrow \text{Enc}(\text{mpk}, x^\beta)$	$\text{ct} \leftarrow \text{Enc}(\text{mpk}, x^\beta)$
$\alpha \leftarrow \mathcal{A}_2^{\text{KeyGen}(\text{mpk}, \text{msk}, \cdot)}(\text{mpk}, \text{ct}, \text{st})$	$\alpha \leftarrow \mathcal{A}_2^{\text{KeyGen}(\text{mpk}, \text{msk}, \cdot), \text{QDec}(\cdot, \cdot)}(\text{mpk}, \text{ct}, \text{st})$
Output: α	Output: α

Fig. 1: IND-CPA and IND-CCA security for functional encryption. The decryption oracle $\text{QDec}(f, \text{ct})$ in the CCA game first generates the secret key $\text{sk}_f = \text{KeyGen}(\text{mpk}, \text{msk}, f)$ and outputs $\text{Dec}(\text{mpk}, f, \text{sk}_f, \text{ct})$ for the query (f, ct) .

Beside the game based security definitions, we also recap a simulation based definition, composable functional encryption (CFE), introduced by Matt and Maurer in [53]. The notion of composable functional encryption (CFE) security.

Definition 2.5 (Composable Functional Encryption Security). Let FE be a functional encryption scheme, $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ a function family, define the experiments $\text{Real}^{\text{FE}}(1^\lambda, \mathcal{A})$ and $\text{Ideal}^{\text{FE}}(1^\lambda, \mathcal{A}, \mathcal{S})$ with a PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ and a PPT simulator $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3)$ respectively in Fig. 2, where the oracle \mathcal{O} is defined as

$$\mathcal{O}(f, x_1, \dots, x_{\ell-1})[[s]] := \mathcal{S}_2(f, f(x_1), \dots, f(x_{\ell-2}))[[s]] .$$

$\text{Real}^{\text{FE}}(1^\lambda, \mathcal{A})$	$\text{Ideal}^{\text{FE}}(1^\lambda, \mathcal{A}, \mathcal{S})$
$(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$	$(\text{mpk}, s) \leftarrow \mathcal{S}_1(1^\lambda)$
$(\ell, \tau) \leftarrow (0, 0)$	$(\ell, \tau) \leftarrow (0, 0)$
Repeat	Repeat
$\ell \leftarrow \ell + 1$	$\ell \leftarrow \ell + 1$
$x_\ell \leftarrow \mathcal{A}_1^{\text{KeyGen}(\text{mpk}, \text{msk}, \cdot)}(\text{mpk})[[\tau]]$	$x_\ell \leftarrow \mathcal{A}_1^{\mathcal{O}(\cdot, x_1, \dots, x_{\ell-1})[[s]]}(\text{mpk})[[\tau]]$
$\text{ct}_\ell \leftarrow \text{Enc}(\text{mpk}, x_\ell)$	$(f_1, \dots, f_q) \leftarrow \text{queries by } \mathcal{A}_1$
$t \leftarrow \mathcal{A}_2(\text{ct}_\ell)[[\tau]]$	$\text{ct}_\ell \leftarrow \mathcal{S}_3(f_0(x_\ell), \dots, f_q(x_\ell))[[s]]$
Until $t = \text{true}$	Until $t = \text{true}$
Output: τ	Output: τ

Fig. 2: CFE security definition

The advantage of the experiments is defined by:

$$\text{Adv}_{\text{FE}, \mathcal{A}, \mathcal{S}}^{\mathcal{D}, \text{CFE}}(\lambda) = |\Pr[\mathcal{D}(\text{Real}^{\text{FE}}(1^\lambda, \mathcal{A})) = 1] - \Pr[\mathcal{D}(\text{Ideal}^{\text{FE}}(1^\lambda, \mathcal{A}, \mathcal{S})) = 1]| ,$$

where \mathcal{D} is a PPT distinguisher.

A functional encryption scheme FE is CFE secure, if there exists a PPT simulator \mathcal{S} , such that for any PPT distinguisher \mathcal{D} it holds that $\text{Adv}_{\text{FE}, \mathcal{A}, \mathcal{S}}^{\mathcal{D}, \text{CFE}}(\lambda) \leq \text{negl}(\lambda)$ for any PPT adversary \mathcal{A} , where $\text{negl}(\cdot)$ is a negligible function.

Remark 2.6 (On the leakage function). As already noted in [53], the leakage function is a modeling artifact specific to the confidentiality definitions: the information captured by f_0 models the general leakage that *might* be possible to compute by an adversary by just observing an *honestly generated* ciphertext, for example the length of the underlying plaintext (which some works put in place by default). Because this information is not guaranteed to be computable f_0 does actually not model a real function as opposed to $f_i, i > 0$. As we will see later, our consistency guarantees will only require that the guaranteed functions $f_i, i > 0$ yield consistent results.

2.3 Standard Tools and Assumptions

Now, we recap the definition of a matrix distribution and the Matrix-Diffie-Hellman assumption as introduced in [35]. We begin with the definition for a matrix distribution.

Definition 2.7 (Matrix Distribution). Let $\ell, k \in \mathbb{N}$ with $\ell > k$. We call $\mathcal{D}_{\ell, k}$ a matrix distribution if it outputs matrices in $\mathbb{Z}_p^{\ell \times k}$ of full rank k in polynomial time. We define $\mathcal{D}_k := \mathcal{D}_{k+1, k}$.

We assume, wlog, that the first k rows of $\mathbf{A} \leftarrow \mathcal{D}_k$ form an invertible matrix. The \mathcal{D}_k -Matrix Diffie-Hellman problem is to distinguish the two distributions $([\mathbf{A}], [\mathbf{A}\mathbf{w}])$ and $([\mathbf{A}], [\mathbf{u}])$ where $\mathbf{A} \leftarrow \mathcal{D}_k$, $\mathbf{w} \leftarrow \mathbb{Z}_p^k$ and $\mathbf{u} \leftarrow \mathbb{Z}_p^{k+1}$.

Now, we state the \mathcal{D}_k -Matrix Diffie-Hellman Assumption (\mathcal{D}_k -MDDH).

Definition 2.8 (\mathcal{D}_k -Matrix Diffie-Hellman Assumption (\mathcal{D}_k -MDDH)). Let \mathcal{D}_k be a matrix distribution. The \mathcal{D}_k -Matrix Diffie-Hellman (\mathcal{D}_k -MDDH) assumption holds relative to GGen if for all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\text{GGen}, \mathcal{A}}^{\mathcal{D}_k\text{-MDDH}}(\lambda) := \Pr[\mathcal{A}(\mathcal{G}, [\mathbf{A}], [\mathbf{A}\mathbf{w}]) = 1] - \Pr[\mathcal{A}(\mathcal{G}, [\mathbf{A}], [\mathbf{u}]) = 1] \leq \text{negl}(\lambda) ,$$

where the probability is taken over $\mathcal{G} = (\mathbb{G}, p, g)$, $\mathbf{A} \leftarrow \mathcal{D}_k$, $\mathbf{w} \leftarrow \mathbb{Z}_p^k$, $\mathbf{u} \leftarrow \mathbb{Z}_p^{k+1}$ and the coin tosses of adversary \mathcal{A} .

2.4 Inner-product Functionality Classes

In this work, we consider two different types of inner-product functionalities as defined in [4]:

Inner Product over \mathbb{Z}_P . Let $\mathcal{F} = \{\mathcal{F}_{P_\lambda}^m\}_{\lambda \in \mathbb{N}}$ be a family (indexed by λ) of sets $\mathcal{F}_{P_\lambda}^m$, where P_λ is a modulus of length λ . Omitting the index λ , the set $\mathcal{F}_P^m = \{f_{\mathbf{y}} : \mathbb{Z}_P^m \rightarrow \mathbb{Z}_P, \text{ for } \mathbf{y} \in \mathbb{Z}_P^m\}$ where

$$f_{\mathbf{y}}(\mathbf{x}) = \langle \mathbf{x}, \mathbf{y} \rangle \bmod P$$

defines the inner-product operation over \mathbb{Z}_P .

Bounded-Norm Inner Product over \mathbb{Z} . Let $\mathcal{F} = \{\mathcal{F}^{m, X_\lambda, Y_\lambda}\}_{\lambda \in \mathbb{N}}$ be a family (indexed by λ) of sets $\mathcal{F}^{m, X_\lambda, Y_\lambda}$. Omitting the index λ , the set $\mathcal{F}^{m, X, Y} = \{f_{\mathbf{y}} : \mathbb{Z}_X^m \rightarrow \mathbb{Z}, \text{ with } \mathbf{y} \in \mathbb{Z}_Y^m\}$, where $\mathbb{Z}_X^m := \{\mathbf{x} \in \mathbb{Z}^m, \text{ with } \|\mathbf{x}\|_\infty < X\}$, $\mathbb{Z}_Y^m := \{\mathbf{y} \in \mathbb{Z}^m, \text{ with } \|\mathbf{y}\|_\infty < Y\}$ and $f_{\mathbf{y}}(\mathbf{x}) = \langle \mathbf{x}, \mathbf{y} \rangle$ defines the bounded-norm inner-product over \mathbb{Z} .

2.5 Non-interactive Proofs

Now, we recapture the definition of non-interactive zero knowledge (NIZK) proofs [18,37,41] and non-interactive witness indistinguishable (NIWI) proofs [14,21,47].

Definition 2.9 (Non-Interactive Zero-Knowledge Proofs). *Let R be an NP Relation and consider the language $L = \{x \mid \exists w \text{ with } (x, w) \in R\}$ (where x is called a statement or instance). A non-interactive zero-knowledge proof (NIZK) for the relation R is a triple of PPT algorithms $\text{NIZK} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$:*

$\text{NIZK.Setup}(1^\lambda)$: Takes as input a security parameter λ and outputs the common reference string CRS.

$\text{NIZK.Prove}(\text{CRS}, x, w)$: Takes as input the common reference string CRS, a statement x and a witness w , and outputs a proof π .

$\text{NIZK.Verify}(\text{CRS}, x, \pi)$: Takes as input the common reference string CRS, a statement x and a proof π , and outputs 0 or 1.

A system NIZK is complete, if (for all $\lambda \in \mathbb{N}$), for all CRS in the support of $\text{Setup}(1^\lambda)$ and all statement-witness pairs in the relation $(x, w) \in R$, it holds that

$$\Pr[\text{NIZK.Verify}(\text{CRS}, x, \text{NIZK.Prove}(\text{CRS}, x, w)) = 1] = 1.$$

Besides completeness, a NIZK system also fulfills the notion of soundness and zero-knowledge, which we introduce in the following two definitions:

Definition 2.10 (Soundness). *Given a proof system $\text{NIZK} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ for a relation R and the corresponding language L , we define the soundness advantage of an adversary \mathcal{A} as the probability:*

$$\text{Adv}_{\text{NIZK}, \mathcal{A}}^{\text{Sound}}(\lambda) := \Pr[\text{CRS} \leftarrow \text{NIZK.Setup}(1^\lambda); (x, \pi) \leftarrow \mathcal{A}(\text{CRS}) : \text{NIZK.Verify}(\text{CRS}, x, \pi) = 1 \wedge x \notin L].$$

A NIZK proof system is called perfectly sound if $\text{Adv}_{\text{NIZK}, \mathcal{A}}^{\text{Sound}}(\lambda) = 0$ for all algorithms \mathcal{A} , and computationally sound, if $\text{Adv}_{\text{NIZK}, \mathcal{A}}^{\text{Sound}}(\lambda) \leq \text{negl}(\lambda)$ for all PPT algorithms \mathcal{A} .

Definition 2.11 (Zero-Knowledge). *Let $\text{NIZK} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ be a NIZK proof system for a relation R and the corresponding language L , $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2)$ a pair of algorithms (the simulator), with $\mathcal{S}'(\text{CRS}, \tau, x, w) = \mathcal{S}_2(\text{CRS}, \tau, x)$ for $(x, w) \in R$, and $\mathcal{S}'(\text{CRS}, \tau, x, w) = \text{failure}$ for $(x, w) \notin R$. For $\beta \in \{0, 1\}$, we define the experiment $\text{ZK}_\beta^{\text{NIZK}}(1^\lambda, \mathcal{A})$ in Fig. 3. The associated advantage of an adversary \mathcal{A} is defined as*

$$\text{Adv}_{\text{NIZK}, \mathcal{A}, \mathcal{S}}^{\text{ZK}}(\lambda) := |\Pr[\text{ZK}_0^{\text{NIZK}}(1^\lambda, \mathcal{A}, \mathcal{S}) = 1] - \Pr[\text{ZK}_1^{\text{NIZK}}(1^\lambda, \mathcal{A}, \mathcal{S}) = 1]|.$$

A NIZK proof system NIZK is called perfect zero-knowledge, with respect to a simulator $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2)$, if $\text{Adv}_{\text{NIZK}, \mathcal{A}, \mathcal{S}}^{\text{ZK}}(\lambda) = 0$ for all algorithms \mathcal{A} , and computationally zero-knowledge, if $\text{Adv}_{\text{NIZK}, \mathcal{A}, \mathcal{S}}^{\text{ZK}}(\lambda) \leq \text{negl}(\lambda)$ for all PPT algorithms \mathcal{A} .

$\mathbf{ZK}_0^{\text{NIZK}}(1^\lambda, \mathcal{A}, \mathcal{S})$	$\mathbf{ZK}_1^{\text{NIZK}}(1^\lambda, \mathcal{A}, \mathcal{S})$
$\text{CRS} \leftarrow \text{NIZK.Setup}(1^\lambda)$	$(\text{CRS}, \tau) \leftarrow \mathcal{S}_1(1^\lambda)$
$\alpha \leftarrow \mathcal{A}^{\text{NIZK.Prove}(\text{CRS}, \cdot, \cdot)}(\text{CRS})$	$\alpha \leftarrow \mathcal{A}^{\mathcal{S}'(\text{CRS}, \tau, \cdot)}(\text{CRS})$
Output: α	Output: α

Fig. 3: Zero-knowledge property of a NIZK proof system.

Furthermore, we say that a NIZK is one-time simulation-sound [61], if the following holds.

Definition 2.12 (One-Time Simulation-Soundness). *Given a proof system $\text{NIZK} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ for an NP relation R with corresponding language L and a simulator $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2)$, we define the simulation-soundness advantage of an algorithm \mathcal{A} by*

$$\text{Adv}_{\text{NIZK}, \mathcal{A}, \mathcal{S}}^{\text{Sim-Sound}}(\lambda) := \Pr[(\text{CRS}, \tau) \leftarrow \mathcal{S}_1(1^\lambda); (x, \pi) \leftarrow \mathcal{A}^{\mathcal{S}_2(\text{CRS}, \tau, \cdot)}(\text{CRS}) : (x, \pi) \notin Q \\ \text{and } x \notin L \text{ and } \text{NIZK.Verify}(\text{CRS}, x, \pi) = 1] ,$$

where Q is the set of all (x', π') , such that \mathcal{A} queried x' to its oracle and π' is the matching response.

A NIZK proof system is called one-time simulation sound with respect to the simulator $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2)$, if $\text{Adv}_{\text{NIZK}, \mathcal{A}, \mathcal{S}}^{\text{Sim-Sound}}(\lambda) \leq \text{negl}(\lambda)$ for all PPT algorithms \mathcal{A} that make at most one query to oracle \mathcal{S}_2 .

Besides these pretty standard NIZK proof systems, we also employ (CRS free) non-interactive witness indistinguishable proofs [14, 21, 47]

Definition 2.13 (Non-Interactive Witness-Indistinguishable Proofs). *Let R be an NP Relation and consider the language $L = \{x \mid \exists w \text{ with } (x, w) \in R\}$ (where x is called a statement or instance). A non-interactive witness-indistinguishable proof (NIWI) for the relation R is a tuple of PPT algorithms $\text{NIWI} = (\text{NIWI.Prove}, \text{NIWI.Verify})$:*

$\text{NIWI.Prove}(1^\lambda, x, w)$: Takes as input the unary representation of the security parameter λ , a statement x and a witness w , and outputs a proof π .

$\text{NIWI.Verify}(1^\lambda, x, \pi)$: Takes as input the unary representation of the security parameter λ , a statement x and a proof π , and outputs 0 or 1.

A system NIWI is complete, if for all statement-witness pairs in the relation $(x, w) \in R$, it holds that

$$\Pr[\text{NIWI.Verify}(1^\lambda, x, \text{NIWI.Prove}(1^\lambda, x, w)) = 1] = 1.$$

A NIWI proof system fulfills additional properties besides completeness, namely soundness and witness-indistinguishability.

Definition 2.14 (Soundness). *Let $\text{NIWI} = (\text{NIWI.Prove}, \text{NIWI.Verify})$ be a NIWI proof system for a relation R and the corresponding language L . We define the advantage of an adversary \mathcal{A} as the following probability:*

$$\text{Adv}_{\text{NIWI}, \mathcal{A}}^{\text{Sound}}(\lambda) := \Pr[(x, \pi) \leftarrow \mathcal{A}(1^\lambda) : \text{NIWI.Verify}(1^\lambda, x, \pi) = 1 \wedge x \notin L].$$

A NIWI proof system NIWI is called perfectly sound if $\text{Adv}_{\text{NIWI}, \mathcal{A}}^{\text{Sound}}(\lambda) = 0$ for all algorithms \mathcal{A} , and computationally sound, if $\text{Adv}_{\text{NIWI}, \mathcal{A}}^{\text{Sound}}(\lambda) \leq \text{negl}(\lambda)$ for all PPT algorithms \mathcal{A} .

$\mathbf{WI}_\beta^{\text{NIWI}}(1^\lambda, \mathcal{A})$ (for relation R)
$(x, w_1, w_2, \text{st}) \leftarrow \mathcal{A}_1(1^\lambda)$
$\pi \leftarrow \text{NIWI.Prove}(1^\lambda, x, w_\beta)$
$\alpha \leftarrow \mathcal{A}_2(\pi, \text{st})$
Output: $\alpha \wedge (x, w_1) \in R \wedge (x, w_2) \in R$

Fig. 4: Witness-indistinguishability of a NIWI proof system. The output condition enforces the use of valid instance witness pairs.

Definition 2.15 (Witness-Indistinguishability). Let $\text{NIWI} = (\text{NIWI.Prove}, \text{NIWI.Verify})$ be a NIWI proof system for a relation R and the corresponding language L . For $\beta \in \{0, 1\}$, we define the experiment $\text{WI}_\beta^{\text{NIWI}}(1^\lambda, \mathcal{A})$ in Fig. 4. The associated advantage of an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ is defined as

$$\text{Adv}_{\text{NIWI}, \mathcal{A}, \mathcal{S}}^{\text{WI}}(\lambda) := |\Pr[\text{WI}_0^{\text{NIWI}}(1^\lambda, \mathcal{A}) = 1] - \Pr[\text{WI}_1^{\text{NIWI}}(1^\lambda, \mathcal{A}) = 1]|.$$

A NIWI proof system is called *witness-indistinguishable*, if $\text{Adv}_{\text{NIWI}, \mathcal{A}}^{\text{WI}}(\lambda) = 0$ for all algorithms $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, and *computationally witness-indistinguishable*, if $\text{Adv}_{\text{NIWI}, \mathcal{A}}^{\text{WI}}(\lambda) \leq \text{negl}(\lambda)$ for all PPT algorithms $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$.

As already described in [12], the construction in [47] relies on the decisional linear (DLIN) assumption and provides perfectly sound non-interactive witness indistinguishability. In [14], the authors rely on a complexity theoretic assumption and also present (less efficient) perfectly sound proofs. For the construction in [21], the authors rely on one-way permutations and indistinguishability obfuscation.

2.6 Verifiable Functional Encryption

Now, we recap the definition of verifiable functional encryption as stated in [12].

Definition 2.16 (Verifiable Functional Encryption). A verifiable functional encryption scheme $\text{VFE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{VerifyCT}, \text{VerifySK})$ extends a functional encryption scheme $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ by two algorithms VerifyCT and VerifySK which have the following behavior:

$\text{VerifyCT}(\text{mpk}, \text{ct})$: Takes as input the master public key mpk and a ciphertext ct and outputs 1 if the ciphertext ct was correctly generated using the master public key mpk for some message x .

$\text{VerifySK}(\text{mpk}, f, \text{sk})$: Takes as input the master public key mpk , a function f and a functional key sk and outputs 1 if the functional key sk was correctly generated as a functional key for the function f .

Beside the correctness and security definition, a verifiable functional encryption scheme also needs to fulfill verifiability:

Definition 2.17 (Verifiability). A verifiable functional encryption scheme VFE for \mathcal{F} is verifiable if, for all $\text{mpk} \in \{0, 1\}^*$, for all $\text{ct} \in \{0, 1\}^*$, there exists $x \in \mathcal{X}$ such that for all $f \in \mathcal{F}$ and $\text{sk} \in \{0, 1\}^*$, the following implication holds:

$$\text{If } \text{VerifyCT}(\text{mpk}, \text{ct}) = 1 \text{ and } \text{VerifySK}(\text{mpk}, f, \text{sk}) = 1$$

then

$$\Pr[\text{Dec}(\text{mpk}, f, \text{sk}, \text{ct}) = f(x)] = 1$$

2.7 Universal Composability

The necessary preliminaries for the UC analysis can be found in Appendix C.

3 Consistency for Functional Encryption Schemes

In this section, we formally define the notion of consistency for functional encryption schemes. We first put forth a systematic treatment of consistency where an entity in a system is considered potentially malicious.

Consistency with respect to various corruption sets. Recall that there are three distinct tasks in functional encryption: parameter/key generation, encryption and decryption which, according to [53], yield three corresponding entities in a system: the input provider, the setup/key manager, and the decryptor. Consistency must be seen as a guarantee that an honest decryptor can rely on even in the presence of malicious other entities. In contrast, confidentiality (in the sense of CPA/CCA or CFE) is a guarantee that an honest input provider relies on against a potentially dishonest decryptor (in the presence of honestly generated setup and keys). We summarize these combinations in Table 1. We remark that aside from the informal justification that the games represent what we intend to capture, we cross-check the games against a constructive and composable model in Section 7 to show that our consistency notions are able to realize the intended *idealized UC-functionality* for FE.

In the following sections, we introduce different consistency notions, where each notion corresponds to a different corruption set of untrusted entities. The different types of consistency are: input consistency (in-CONS), strong input consistency (st-in-CONS), and setup consistency (set-CONS).

Notions	Entities		
	Input Provider	Setup+Key Generator	Decryptor
Correctness	Honest	Honest	Honest
in-CONS	Corrupted	Honest	Honest
set-CONS	Honest	Corrupted	Honest
st-in-CONS	Corrupted	Corrupted	Honest
Confidentiality	Honest	Honest	Corrupted

Table 1: The different consistency notions and the corrupted entities

In the rest of this section, whenever we refer to a function f , or a functionality class \mathcal{F} , we implicitly mean the induced function extension as defined in Definition 2.2.

3.1 Consistency with a dishonest Input Provider

An input consistency attack entails the malicious generation of a ciphertext ct , and the honest generation of several non-trivial functional keys $sk_{f_1}, \dots, sk_{f_n}$ that interpret the ciphertext ct in an inconsistent way. Informally, we call a ciphertext consistent, if there exists a plaintext x that can explain the decryption of the ciphertext ct under the different functional keys $sk_{f_1}, \dots, sk_{f_n}$. While consistency should express that a set of output values does have a common explanation, it is desirable to additionally capture conditions under which this explanation is actually unique and determinable. To be able to express this additional property that the scheme can have, the security game may include an efficient algorithm, **Extract**, that given a ciphertext ct and a set of functional decryption keys (with their corresponding functions) returns either **unknown** or a bit string $x \in \mathcal{X}$: in the former case, the game formalizes the basic consistency guarantees, and in the latter case, the game formalizes the stronger guarantee that the value x must be part of the preimage and hence ct is committed to contain x ⁴.

Definition 3.1 (Predicate $Q_\lambda^{\mathcal{F}_\lambda}$). Given a functionality class \mathcal{F}_λ we denote by $Q_\lambda^{\mathcal{F}_\lambda}$ a monotone⁵ predicate on function sets $\mathcal{P}(\mathcal{F}_\lambda) \rightarrow \{0, 1\}$.

⁴ The exact role of **Extract** will be important once we discuss consistency in simulation-based frameworks where a simulator often needs to extract the plaintext from the ciphertext (having simulated the secret keys).

⁵ Meaning that if $Q_\lambda^{\mathcal{F}_\lambda}(F) = 1$ and $F \subseteq F'$ then $Q_\lambda^{\mathcal{F}_\lambda}(F') = 1$.

We are ready to formally state input-consistency:

<pre> in-CONS^{FE}(1^λ, \mathcal{A}, Extract) ----- (msk, mpk) ← Setup(1^λ) ct ← $\mathcal{A}^{\text{KeyGen}(\text{mpk}, \text{msk}, \cdot)}$(1^λ, mpk) Let $F := \{(\text{sk}_{f_i}, f_i)\}_{i \in [n]}$ be the set of key-function pairs obtained by \mathcal{A}. If $n < 1$ then output 0 $x := \text{Extract}(\text{mpk}, \text{ct}, F)$ $y_i := \text{Dec}(\text{mpk}, f_i, \text{sk}_{f_i}, \text{ct})$, for all $i \in [n]$ If $x \neq \text{unknown}$ If $\bigcap_{i \in [n]} f_i^{-1}(y_i) \neq \{x\}$ Output 1 Output 0 If $x = \text{unknown}$ If $\bigcap_{i \in [n]} f_i^{-1}(y_i) = \emptyset$ Output 1 Output 0 </pre>
--

Fig. 5: Input Consistency Definition.

Definition 3.2 (Input Consistency). Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be a functional encryption scheme, then we define the experiment $\text{in-CONS}^{\text{FE}}$ in Fig. 5. The scheme FE is input consistent, if there exists an extract algorithm Extract with the format as specified above such that for any polynomial-time adversary \mathcal{A} , there exists a negligible function negl such that:

$$\Pr[\text{in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract})] \leq \text{negl}(\lambda) .$$

In addition, we say that the algorithm Extract satisfies $\mathbb{Q}_\lambda^{\mathcal{F}^\lambda}$ -input-extractability, if for any PPT adversary \mathcal{A} in the experiment $\text{in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract})$, it holds that:

$$\Pr[\mathbb{Q}_\lambda^{\mathcal{F}^\lambda}(\{f_i \mid (\cdot, f_i) \in F\}) = 1 \implies x \neq \text{unknown}] = 1,$$

where the probability is taken over the random coins of the adversary and the scheme, and where F and x are defined by the experiment.

The additional input of an extract algorithm Extract that can be provided to the consistency game allows us to distinguish between three different cases:

1. **The algorithm Extract is trivial, i.e., always outputs unknown.** This is the most basic and weakest form of consistency. In this case extractability is not taken into account and it is just checked that a ciphertext is explainable under different functional keys.
2. **The algorithm Extract is not trivial.** This is a stronger form of consistency, where the extract algorithm is able to extract the underlying message of some ciphertexts. We require that whenever the extract algorithm extract a message x that this must be the only possible message that explains the decryption behavior of ct under the different functional keys $\text{sk}_{f_1}, \dots, \text{sk}_{f_n}$ (together with their corresponding functions). In other words, this ciphertext ct is committed to a value x .

3. The algorithm Extract is not trivial and satisfies $\mathcal{Q}_\lambda^{\mathcal{F}_\lambda}$ -input-extractability. This is the strongest form of consistency, where extraction is possible on every ciphertext ct given a $\mathcal{Q}_\lambda^{\mathcal{F}_\lambda}$ -input-extractable function set F . This means we do not only know that any ciphertext is committed to an underlying plaintext as a property of the functionality class, but also that $\mathcal{Q}_\lambda^{\mathcal{F}_\lambda}$ captures when this must be the case.⁶

We note in passing that $3 \Rightarrow 2 \Rightarrow 1$.

Discussion. It is worth mentioning how the game reflects (in-)consistency: After the adversary \mathcal{A} has obtained the functional keys $\text{sk}_{f_1}, \dots, \text{sk}_{f_n}$ for all the functions $\{f_1, \dots, f_n\}$ it has asked to the key generation oracle $\text{KeyGen}(\text{msk}, \cdot)$, it outputs a ciphertext ct trying to win the game. The adversary has two chances to do so: First, the challenger uses the generated functional keys together with the corresponding functions $\{\text{sk}_{f_i}, f_i\}_{i \in [n]}$ and tries to extract a plaintext x , out of the ciphertext ct , by executing the Extract algorithm. If this yields a valid message x , it checks if the functional behavior of this message is the same as the functional behavior on the ciphertext. In more detail, the challenger checks if $\text{Dec}(\text{sk}_{f_i}, \text{ct}) = f_i(x)$ for all $i \in [n]$. If this check does not hold, the adversary has generated a malicious ciphertext that produces inconsistent output values, i.e., outputs that are not explained by the extracted value x . In addition, if the scheme is to satisfy \mathcal{Q} -input-extractability, then not only such a value x must exist, but x must be determined by this point if the set of functions queried by \mathcal{A} satisfies \mathcal{Q} .

Since not every set F of functions allows extraction, there is a second phase, where “plain consistency” is checked: If the algorithm Extract does not return a value x , the challenger checks whether all produced outputs of the decryption algorithm can have a common explanation: The challenger checks if there exists a set of messages that explains the functional decryption behavior of ct under the different functional keys $\{\text{sk}_{f_i}\}_{i \in [n]}$. Formally, it computes the intersection $\bigcap_{i \in [n]} f_i^{-1}(\text{Dec}(\text{mpk}, f_i, \text{sk}_{f_i}, \text{ct}))$ and if it is empty, there is no explanation for the decryption behavior of ct , which means that the adversary has caused an inconsistency and wins the game. Note that the plain consistency check does not imply that ct is a commitment to a particular value. Nevertheless, consistency can be evaluated.

Also note that the intersection is well defined but not always efficiently computable, for example when the f_i 's are one-way functions. Whether a restriction of the function class w.r.t. efficiently computable preimages is necessary depends on the bigger construction in which the FE scheme is employed—and in particular on their reduction proof.⁷ On the other hand, when used as an assumption in a proof, then the efficiency restriction is a simple way to make the assumption falsifiable [54]. We make use of such a restriction for our UC proof in Section 7.

It is furthermore interesting to see how the special symbols \perp and \diamond obtain the intended semantic by this game: If one of the decryption algorithm invocation outputs \diamond at any time in the game, the adversary wins because the preimage of \diamond under every function is the emptyset, i.e. $f^{-1}(\diamond) = \emptyset$ (see Definition 2.2), which results in an overall empty intersection; in particular there exists no x in the message space $x \in \mathcal{X} \cup \{\perp\}$ such that $f(x) = \diamond$ due to the definition of the function extension (Definition 2.2), which makes the adversary win the game. This captures that when the public parameters and the functional keys are honestly generated, then the decryption algorithm is not allowed to output \diamond (recall that the symbol is used to indicate an incorrect key).

Analogously, if one of the decryption algorithm invocations outputs \perp and another decryption algorithm invocation outputs a value $y_i \neq \perp$ then the adversary wins the game, as the intersection must be empty since the preimage of \perp is \perp which is not equal to the preimage of y_i (Definition 2.2). This captures that the corresponding ciphertext cannot be honestly generated, due to the disagreement of the keys on the validity of the ciphertext.

Remark 3.3 (On the leakage function). As noted earlier, we deliberately ignore the leakage function f_0 when defining consistency requirements, since we perceive f_0 , as already noted in [53], as a modeling artifact specific

⁶ A natural way to imagine this is a functionality class that has a set of basis functions that spawn the entire functionality class. Finding such a basis might not be easy in general.

⁷ Note that similar thoughts apply, e.g., to extractor games in interactive zero-knowledge proofs where the experiment need not be bounded by a polynomial.

<p style="margin: 0;">st-in-CONS^{FE}($1^\lambda, \mathcal{A}, \text{Extract}$)</p> <hr style="border: 0.5px solid black;"/> <p style="margin: 0;">($\text{mpk}, \text{ct}_1, \text{ct}_2, \{(\text{sk}_j, f_j)\}_{j \in [n]}\} \leftarrow \mathcal{A}(1^\lambda)$ (Assume $\text{sk}_j \neq \varepsilon$)</p> <p style="margin: 0;">$y_{i,j} := \text{Dec}(\text{mpk}, f_j, \text{sk}_j, \text{ct}_i)$, for all $j \in [n], i \in \{1, 2\}$</p> <p style="margin: 0;">If $y_{1,j} = \diamond \wedge y_{2,j} \neq \diamond$ or $y_{1,j} \neq \diamond \wedge y_{2,j} = \diamond$ for any $j \in [n]$</p> <p style="margin: 0; padding-left: 20px;">Output 1</p> <p style="margin: 0;">Let $F := \{(\text{sk}_j, f_j)\}_{j \in [n] \wedge (y_{1,j} \neq \diamond \vee y_{2,j} \neq \diamond)}$</p> <p style="margin: 0;">If F is empty then output 0</p> <p style="margin: 0;">$x_i := \text{Extract}(\text{mpk}, \text{ct}_i, F)$, for all $i \in \{1, 2\}$</p> <p style="margin: 0;">For each $i \in \{1, 2\}$ do:</p> <p style="margin: 0; padding-left: 20px;">If $x_i \neq \text{unknown}$</p> <p style="margin: 0; padding-left: 40px;">If $\bigcap_{j \in [n], (\cdot, f_j) \in F} f_j^{-1}(y_{i,j}) \neq \{x_i\}$</p> <p style="margin: 0; padding-left: 60px;">Output 1</p> <p style="margin: 0; padding-left: 20px;">If $x_i = \text{unknown}$</p> <p style="margin: 0; padding-left: 40px;">If $\bigcap_{j \in [n], (\cdot, f_j) \in F} f_j^{-1}(y_{i,j}) = \emptyset$</p> <p style="margin: 0; padding-left: 60px;">Output 1</p> <p style="margin: 0;">Output 0</p>

Fig. 6: Strong Input Consistency Definition.

to the confidentiality definitions that we do not need to port to our new definition: the information captured by f_0 models the general leakage that *might* be possible to compute by an adversary by just observing an *honestly generated* ciphertext. However, it seems unreasonable to assume that this must be guaranteed to be computable. For instance, in the case of standard encryption schemes, computing the length of the plaintext is not guaranteed by the scheme, but the definition does formally not require that this information must be hidden. The distinction will appear more clearly in the UC treatment of consistency, where f_0 is just treated as leakage to the adversary.

3.2 Consistency with a dishonest Input Provider and Key Generator

We turn our attention to a stronger coalition against an honest decryptor, which is the setting in which both, the input provider and the parameter/key generation entities are dishonest. The adversary outputs a malicious master public key mpk , two ciphertexts ct_1, ct_2 and a set of functional keys $\{\text{sk}_{f_i}\}_{i \in [n]}$. The attacker's goal is that the latter keys interpret the ciphertexts ct_1, ct_2 in an inconsistent way. Contrary to input consistency, the game considers two ciphertexts. This allows us to capture consistency of key validity as modeled by \diamond in the 3rd line of the game. While technically redundant due to the disjunction in F , we chose to highlight this important property in the game. As, with input consistency, an adversary breaks consistency, if there exists no plaintext x , for at least one of the challenge ciphertexts, that can explain the decryption of some ciphertext ct under the different functional keys $\text{sk}_{f_1}, \dots, \text{sk}_{f_n}$.

Similar to input consistency, we capture the fact that the ciphertext might be committing given a specific set of keys, and again introduce an algorithm Extract in the security to capture this if needed.⁸ More formally:

Definition 3.4 (Strong Input-Consistency). Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be a functional encryption scheme, then we define the experiment $\text{st-in-CONS}^{\text{FE}}$ in Fig. 6. The scheme FE is strong input consistent,

⁸ As for input-consistency, setting $\text{Extract} := \text{unknown}$ in the experiment is an option to simply capture plain consistency in this setting.

if there exists an extract algorithm Extract with the format as specified above, such that for any polynomial-time adversary \mathcal{A} , there exists a negligible function negl such that:

$$\Pr[\text{st-in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract})] \leq \text{negl}(\lambda) .$$

In addition, we say that the algorithm Extract satisfies strong $\mathcal{Q}^{\mathcal{F}_\lambda}$ -extractability, if for any PPT adversary \mathcal{A} in the experiment $\text{st-in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract})$, it holds that:

$$\Pr[\mathcal{Q}^{\mathcal{F}_\lambda}(\{f_j \mid (\cdot, f_j) \in F\}) = 1 \implies x_i \neq \text{unknown}] = 1,$$

where the probability is taken over the random coins of the adversary and the scheme, and where F and x_i are defined by the experiment.

Discussion. The consistency experiment above is motivated by augmenting the attack capabilities of the previous one. Here, the adversary can output a master public key mpk , ciphertexts $\{\text{ct}_i\}_{i \in [2]}$ and a set of functional keys $\{\text{sk}_j, f_j\}_{j \in [n]}$ (again, note that we do not give any guarantees for f_0 and the empty key).

Compared to the weaker form in the previous section, however, not all keys will be deemed valid, and hence the set F is defined as the set of those key-function pairs (sk_j, f_j) that, for at least one ciphertext ct_i yield a decryption $y_{i,j} \neq \diamond$. Only keys in F can provoke a consistency breach. This thereby again assigns the intended meaning that obtaining \diamond indicates that the decryption algorithm deems a key invalid (and therefore, as we detail below, does only provoke an inconsistency if not all decryptions w.r.t. this key yield \diamond).

Using F , the challenger again tries to extract a preimage x_i for each ciphertext ct_i and the consistency checks proceed similar to input-consistency: If x_i was recovered, the game checks if the decryption of the ciphertext ct_i under the different functional keys $\{\text{sk}_{f_j}\}_{j \in [n]}$ is the same as the function applied on x_i , i.e. $\text{Dec}(\text{sk}_{f_j}, \text{ct}_i) = f_j(x_i)$ for all $(\text{sk}_j, f_j) \in F$ (expressed using the intersection-notation). If this is not the case, the adversary has generated a ciphertext, for which the behavior under the different functional keys cannot be explained and hence would win the game.

If an x_i was not recovered, the existence of a common explanation is checked, i.e. whether there exists a message in the intersection of the preimages under the different functions $\{f_j\}_{j \in [n], (\cdot, f_j) \in F}$. If the intersection is empty, the adversary has generated a ciphertext ct_i with a decryption behavior that cannot be explained.

Again, the symbols \diamond and \perp deserve a special observation. As before, if a key yields consistently the “decryption” \diamond , this key is detected as invalid; otherwise, if for some ciphertext ct_i we have that exactly one decryption $y_{i,j} \neq \diamond$, the performed intersection check must yield the empty set by Definition 2.2. In other words, this behavior captures our intention that a key’s validity cannot vary depending on what is decrypted, as the decryptor in one case believes it is holding a valid key for the function f .

On the other hand, if a ciphertext is deemed invalid, i.e., $y_{i,j} = \perp$, then all functions in F must consistently declare this ciphertext as invalid, as otherwise, the intersection will again be empty by the definition of symbol \perp in Definition 2.2. Furthermore, if Extract returns a value, it must be \perp as well.

Finally, we note again that if a set of functions is strongly \mathcal{Q} -extractable⁹ for some predicate \mathcal{Q} , then this means that given the set of valid functional keys and a ciphertext that the ciphertext is committing to a specific value $x_i \in \mathcal{X} \cup \{\perp\}$.

3.3 Consistency with a dishonest Parameter/Key Generator

We define consistency in the case of an untrustworthy parameter/key generator. This form is called setup-consistency. At first sight it seems as if this “intermediate” notion is already covered by the previous sections but, surprisingly, it is not. Setup consistency captures the important case where an authority might tamper with the system’s parameters and hence this notion captures retaining consistent in the presence of subversion attacks [11, 16]. We formalize consistency attacks (i.e., the capabilities of an adversary) by letting the adversary define inputs (out of which honest ciphertexts are generated), and where the master public key mpk and

⁹ We just write \mathcal{Q} -extractable because it is usually clear from the context.

a functional key sk are produced entirely by the attacker. Note that we allow the attacker to specify two master public-keys (one for the input provider and one for the decryptor)¹⁰. An attack breaks consistency, if the functional key sk together with the function f yields inconsistent output values with respect to the ciphertexts, i.e. the decryption of the ciphertexts under the functional key sk reveals a mismatch with respect to the input values and the declared function f (unless sk is identified as bogus). As for strong input consistency, consistency of key validity as modeled by \diamond is captured using a disjunction, this time in the outer “If” statement. In more detail, we define the following:

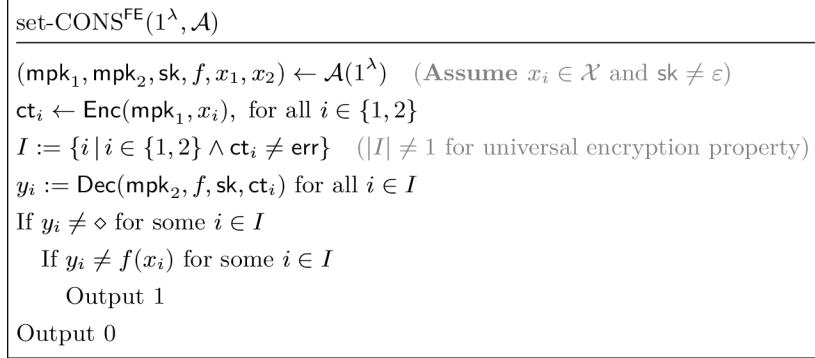


Fig. 7: Setup Consistency Definition.

Definition 3.5 (Setup Consistency). Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be a functional encryption scheme, then we define the experiment $\text{set-CONS}^{\text{FE}}$ in Fig. 7. The scheme FE satisfies setup consistency (or set-CONS for short), if for any polynomial-time adversary \mathcal{A} , there exists a negligible function negl such that:

$$\Pr[\text{set-CONS}^{\text{FE}}(1^\lambda, \mathcal{A})] \leq \text{negl}(\lambda) .$$

In addition, we say that the scheme satisfies the universal encryption property, if in the above experiment, $|I| \in \{0, 2\}$ with overwhelming probability (where the probability is taken over the random coins of the adversary and the scheme and I is defined by the game).

Discussion. It is again instructive to see the nature of consistency attacks that an adversary can mount against a scheme. After the adversary \mathcal{A} outputted two master public keys mpk_1 and mpk_2 , a functional key sk , a function f and two chosen messages x_1 and x_2 , the challenger encrypts the messages under mpk_1 to generate $\text{ct}_i = \text{Enc}(\text{mpk}_1, x_i)$. Now, we are interested in the functional behavior of all valid encryptions that the input provider produces (i.e., that do not return an err symbol upon encryption because of a bogus mpk_1). At this point, a stronger property that we term *universal encryption* could be desirable for applications. Namely, to require that either both encryptions are valid or none is. If the property does not hold, a maliciously generated mpk_1 may only allow for the encryption of a subset of the plaintext space, which might be problematic in some applications. Similar to the extractability property before, this property should be considered as an add-on if needed, but not as lying at the core of a consistency definition. For example, such an all-or-nothing encryption property can be obtained whenever an efficiently computable membership-test for the support of Setup is available for an FE scheme with perfect correctness. We refer to the UC treatment to show what we gain in terms of additional security if insisting on the universal encryption property.

Let us, for concreteness, discuss the case where both encryptions are valid: if both decryption invocations under sk return the special symbol \diamond then the adversary does not win the game (in this case, the key is

¹⁰ Thanks to our UC treatment, we also see the need for this: if there is only one master public-key in the game, then this would imply that one assumes a broadcast channel between the dishonest setup generator and the honest input provider and decryptor.




Fig. 8: Relations between consistency and confidentiality, where the crossed arrows indicate “does not imply” and “&” denotes both properties simultaneously.

deemed bogus). However, if only one of the two outputs \diamond the adversary immediately wins the game (as there can be no value x_i in the domain that yields $f(x_i) = \diamond$ (see Definition 2.2)). Recall that this behavior captures our intention that a key’s validity cannot vary depending on what is decrypted, as the decryptor in one case believes it is holding a valid key for the function f . Now, we consider the case where both decryption attempts yield values $y_i \neq \diamond$. In this case, to fulfill consistency, both of these values must satisfy $f(x_i) = y_i$, otherwise the attacker has broken consistency. If the decryption procedure would output $y_i = \perp$ a security breach happens. In more detail, considering that honestly generated ciphertexts are committed to a real message (otherwise the decryption must be considered as inconsistent), the adversary wins the game. The reason that the adversary wins in this case follows by Definition 2.2 since no message $x_i \neq \perp$ maps to \perp .

Strong robustness against subversion. Looking ahead to Section 7 where we present the justification of the game by showing that it admits the realization of a natural ideal repository with access control, we see that in fact, we must insist that the inputs provided by Alice do functionally match the values that Bob decrypts. Otherwise, the guarantee for honest parties in this setting with subversion of parameters would be too weak, as it would merely enforce consistent decryption—but potentially with respect to a common preimage x' never intended by Alice! This is a form of robustness not implied by strong input consistency or verifiable functional encryption [12]. This shows that a separate notion for the case of subverted setup, namely setup-consistency, is indeed desirable to formalize. In particular, it is necessary for our UC treatment of consistency.

4 Relations (in)between Consistency and Confidentiality

In this section, we formally examine the relationship among the consistency notions and between security and consistency. The main result is depicted in Fig. 8 and shows that consistency is a property independent of the known confidentiality properties. Namely, we prove that Consistency does not imply Confidentiality in Section 4.2, that Confidentiality does not imply consistency in Section 4.3 and that Consistency does not imply a confidentiality lifting in Section 4.4.

If not otherwise quantified, we denote by \mathcal{F} a functionality class, the members of \mathcal{F} by f_i , and refer to the number of functions (not counting the distinguished leakage function f_0) as the size s of the functionality class.

4.1 Relations among the Consistency Notions

Let us first summarize the relations between the notions which can all be seen by simple arguments: strong input consistency implies input consistency since the attack model of input consistency is a strict subset of strong input consistency. Furthermore, since the schemes we present in Section 5 are input consistent but neither strong input consistent nor setup consistent. The only remaining non-implications are that strong input consistency does not imply setup consistency and that setup consistency does not imply strong or normal input consistency. Formally, both are easy to see: one can always take an input or strong-input

consistent scheme and introduce a special master public key mpk' (that has probability zero of being generated by setup) which takes all messages to special ciphertext $\bar{\text{ct}}$ that decrypts to \perp . Such a scheme is obviously not setup consistent but remains consistent because ct decrypts consistently. Along the same lines, one can introduce a new special ciphertext $\bar{\text{ct}}$ in a setup consistent scheme, that decrypts to inconsistent outputs but clearly has probability 0 to be output by the encryption algorithm. This scheme remains setup consistent but is clearly not input consistent.

4.2 Consistency does not imply Confidentiality

To show that consistency does not imply confidentiality, we aim to construct a scheme that satisfies st-in-CONS but is not IND-CPA secure. The scheme is described in Fig. 9. It is easy to see that the scheme described in Fig. 9 does not provide any confidentiality guarantee since the ciphertext reveals the input message. We prove the consistency of the scheme more formally:

Theorem 4.1 (Strong input consistency). *The functional encryption scheme $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ described in Fig. 9 is strongly input consistent for any functionality class $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$. Namely, there exists an algorithm Extract such that for any PPT adversary \mathcal{A} , it holds that:*

$$\Pr[\text{st-in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract})] \leq \text{negl}(\lambda) .$$

<p><u>Setup(1^λ) :</u> Return $(\text{mpk}, \text{msk}) \leftarrow \{0, 1\}^\lambda \times \{0, 1\}^\lambda$</p> <p><u>KeyGen($\text{msk}, f$) :</u> Return $\text{sk}_f = f$</p> <p><u>Enc(mpk, x) :</u> Return $\text{ct} = x$</p> <p><u>Dec($\text{mpk}, f, \text{sk}_f, \text{ct}$) :</u> Parse $\text{ct} = x$ If $x \notin \mathcal{X}_\lambda$ Return \perp If $\text{sk}_f \neq f \vee f \notin \mathcal{F}_\lambda$ Return \diamond Return $f(x)$</p>

Fig. 9: A strongly input consistent FE scheme which is not IND-CPA secure.

Proof. The extract-algorithm in this case can simply return the first argument, as encryption is the identity function on bitstrings (in this sense, the scheme even fulfills Q-extractability for any Q).

Now, we analyze strong input consistency. After the challenger has received a ciphertext ct and some functional keys $\{\text{sk}_{f_i}, f_i\}_{i \in [n]}$ from \mathcal{A} , it executes the extract algorithm $\text{Extract}(\text{ct}, \{f_i\}_{i \in [n]})$ and obtains $\text{ct} = x$ by definition of the encryption algorithm. The challenger then checks if $\text{Dec}(\text{mpk}, f, \text{sk}_{f_i}, \text{ct}) \neq f_i(x)$ for at least one $i \in [n]$. However, due to the definition of the key generation algorithm, (i.e. $\text{KeyGen}(\text{msk}, f) = f$), it holds that $\text{Dec}(\text{mpk}, f, \text{sk}_{f_i}, \text{ct}) = f_i(x)$ for all $i \in [n]$ and the scheme must be consistent. \square

The scheme is trivially setup consistent since the encryptor ignores any setup values and Bob just evaluates the plain functions. Finally, input consistency follows, since it is implied by strong input consistency.

4.3 Confidentiality does not imply Consistency

Next, we prove that the strongest confidentiality notions in use, i.e., IND-CCA and CFE, do not imply consistency with respect to dishonest input provider or parameter generator.

The IND-CCA case. At first glance, the notions of IND-CCA security and input consistency seem to be related. In both games, the scheme must tame the adversaries capabilities of generating malicious ciphertexts. We show however that there is no connection between IND-CCA security and input or setup consistency, by presenting a scheme that is IND-CCA secure, but not input or setup consistent. The scheme is described in Fig. 10, it is based on the brute-force construction of [23, Section 4].

```

Setup( $1^\lambda$ ) :
For  $i = 1, \dots, s$  run  $(pk_i, sk_i) \leftarrow \text{PKE.Setup}(1^\lambda)$ 
 $(mpk, msk) = (\{pk_i\}_{i \in [s]}, \{sk_i\}_{i \in [s]})$ 
KeyGen( $msk, f_i$ ) :
Parse  $msk = \{sk_i\}_{i \in [s]}$ 
Return  $sk_{f_i} = sk_i$ 
Enc( $mpk, x$ ) :
Parse  $mpk = \{pk_i\}_{i \in [s]}$ 
Compute  $ct_i = \text{PKE.Enc}(pk_i, f_i(x)), \forall i \in [s]$ 
Return  $ct = (ct_i)_{i \in [s]}$ 
Dec( $mpk, sk_{f_i}, f_i, ct$ ) :
Parse  $mpk = \{pk_i\}_{i \in [s]}, sk_{f_i} = sk_i, ct := (ct_i)_{i \in [s]}$ 
Return  $y := \text{PKE.Dec}(sk_i, ct_i)$ 

```

Fig. 10: An IND-CCA secure, but not consistent functional encryption scheme.

Theorem 4.2. *Let $\text{PKE} = (\text{PKE.Setup}, \text{PKE.Enc}, \text{PKE.Dec})$ be an IND-CCA secure public-key encryption scheme, then the functional encryption scheme $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ in Fig. 10 is IND-CCA secure for any functionality class \mathcal{F} of polynomial size s (in the security parameter). Namely, for any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B} such that*

$$\text{Adv}_{\text{FE}, \mathcal{A}}^{\text{IND-CCA}}(\lambda) \leq s \cdot \text{Adv}_{\text{PKE}, \mathcal{B}}^{\text{IND-CCA}}(\lambda) .$$

Proof. To prove this statement, we use a hybrid argument over the games G_0, \dots, G_s as defined in Fig. 11. Note that G_0 corresponds to the game $\text{IND-CCA}_0^{\text{FE}}$ and game G_s to the game $\text{IND-CCA}_1^{\text{FE}}$. By using the triangle inequality, we get:

$$\text{Adv}_{\text{FE}', \mathcal{A}}^{\text{IND-CCA}}(\lambda) \leq \sum_{k=1}^s |\text{Win}_{\mathcal{A}}^{G_{0,k-1}} - \text{Win}_{\mathcal{A}}^{G_{0,k}}| .$$

We conclude the proof by showing that for any $k \in [s]$, there exists an adversary \mathcal{B}_k such that:

$$|\text{Win}_{\mathcal{A}}^{G_{0,k-1}} - \text{Win}_{\mathcal{A}}^{G_{0,k}}| \leq \text{Adv}_{\text{PKE}, \mathcal{B}_k}^{\text{IND-CCA}}(\lambda) .$$

The adversary \mathcal{B} of the statement is then defined as the monolithic adversary that first samples $k \leftarrow [s]$ uniformly at random and then runs the code of \mathcal{B}_k .

Game	ct	justification/remark
G_0	$\text{Enc}(\text{mpk}, f_i(x^0))$, for all $i \in [s]$	$G_0 = G_{0,0}$
$G_{0,k}$	$\text{Enc}(\text{mpk}, f_i(x^1))$, for all $i \leq k$ $\text{Enc}(\text{mpk}, f_i(x^0))$, for all $i > k$	IND-CCA of PKE
G_1	$\text{Enc}(\text{mpk}, f_i(x^1))$, for all $i \in [s]$	$G_1 = G_{0,s}$

Fig. 11: Overview of the games to prove the IND-CCA security of the functional encryption scheme described in Fig. 10.

We build an adversary \mathcal{B}_k that simulates $G_{0,k-1+\beta}$ to \mathcal{A} , when interacting with the underlying $\text{IND-CCA}_\beta^{\text{PKE}}$ experiment.

In the first step of the reduction, the adversary \mathcal{B}_k receives the public key pk from the experiment. It sets $\text{pk}_k = \text{pk}$ and generates public key instances $(\text{pk}_i, \text{sk}_i) \leftarrow \text{PKE.Setup}(1^\lambda)$ for all $i \in [s] \setminus \{k\}$, defines the master public key as $\text{mpk} := \{\text{pk}_i\}_{i \in [s]}$ and sends it to \mathcal{A} .

Whenever \mathcal{A} asks for a functional key sk_{f_i} , with $i \in [s] \setminus \{k\}$, \mathcal{B}_k outputs sk_i to \mathcal{A} . If \mathcal{A} asks for the functional key sk_{f_k} the adversary \mathcal{B}_k outputs a random value $\alpha \leftarrow \{0, 1\}$ as its guess. Note that by definition, \mathcal{A} is in this case restricted to submit identical challenge messages w.r.t. the public key pk for which case the behavior of $G_{0,k-1}$ and $G_{0,k}$ are identical (and thus independent of β).

When \mathcal{A} submits its challenge messages (x^0, x^1) , the adversary \mathcal{B}_k computes $\text{ct}_i = \text{PKE.Enc}(\text{pk}_i, f_i(x^1))$ for all $i < k$ and $\text{ct}_i = \text{PKE.Enc}(\text{pk}_i, f_i(x^0))$ for all $i > k$. To generate the ciphertext ct_k , \mathcal{B}_k creates the challenge $(f_k(x^0), f_k(x^1))$ and submits it as its own challenge. \mathcal{B}_k receives $\text{ct}_k = \text{PKE.Enc}(\text{pk}_k, f_k(x^\beta))$ as an answer and creates the ciphertext $\text{ct} := (\text{ct}_i)_{i \in [s]}$, which it sends to \mathcal{A} .

If \mathcal{A} queries the decryption oracle $\text{QDec}(f_i, \text{ct})$ with $i \in [s] \setminus \{k\}$, \mathcal{B}_k computes $f_i(x) = \text{PKE.Dec}(\text{sk}_i, \text{ct})$ and sends $f_i(x)$ to \mathcal{A} . In the case \mathcal{A} queries $\text{QDec}(f_k, \text{ct})$, \mathcal{B}_k forwards ct to its own decryption oracle and sends the reply to \mathcal{A} .

In the last step, the adversary \mathcal{B}_k outputs the same bit β' returned by \mathcal{A} . Since \mathcal{B}_k perfectly emulates $G_{0,k-1+\beta}$ to \mathcal{A} as long as the public key pk_k is not asked, and since for the latter exception case, the advantage of \mathcal{A} is zero, \mathcal{B}_k 's distinguishing advantage in the CCA game is at least the advantage of \mathcal{A} distinguishing systems $G_{0,k-1+\beta}$, for $\beta \in \{0, 1\}$. \square

After showing the IND-CCA security of the scheme, we describe a successful attacker for the input consistency game.

Theorem 4.3. *Let $\text{PKE} = (\text{PKE.Setup}, \text{PKE.Enc}, \text{PKE.Dec})$ be an IND-CCA secure public-key encryption scheme, then the functional encryption scheme $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ in Fig. 10 is not input consistent for a concrete functionality class \mathcal{F} of size $s = 2$ (as described in the proof). Namely, for any algorithm Extract , there exists a PPT adversary \mathcal{A} such that*

$$\Pr[\text{in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract}) = 1] = 1 .$$

Proof. We consider a functionality class that contains two functions ($s = 2$), i.e. $\mathcal{F} = \{f_1, f_2\}$, with $f_1 : \mathcal{X}_\lambda \rightarrow \{0, 1\}$ and $f_2(x) := \overline{f_1(x)}$, where $\overline{\cdot}$ denotes the bit complement. The adversary \mathcal{A} generates a ciphertext $\text{ct} = (\text{ct}_1, \text{ct}_2) = (\text{Enc}(\text{pk}_1, 0), \text{Enc}(\text{pk}_2, 0))$, asks the KeyGen oracle for the two secret keys and sends (ct, f_1, f_2) to the challenger.

We observe that independent of the output of $\text{Extract}(\text{ct}, \{\text{sk}_i, f_i\}_{i \in \{1, 2\}})$, both decryptions will yield $y_i = 0$ as an output. Therefore, we have obtain in any case $f_1^{-1}(0) \cap f_2^{-1}(0) = f_1^{-1}(0) \cap \overline{f_1^{-1}(0)} = \emptyset$, which contradicts the input-consistency requirement. \square

The scheme described in Fig. 10 is also not setup consistent.

Theorem 4.4. *Let $\text{PKE} = (\text{PKE.Setup}, \text{PKE.Enc}, \text{PKE.Dec})$ be an IND-CCA secure public-key encryption scheme, then the functional encryption scheme $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ in Fig. 10 is not setup consistent for a concrete function class \mathcal{F} of size $s = 2$ (as described in the proof). Namely, there exists a PPT adversary \mathcal{A} such that*

$$\Pr[\text{set-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}) = 1] = 1 .$$

Proof. Similarly to the proof of input consistency, we consider a functionality class that contains two functions ($s = 2$), i.e. $\mathcal{F} = \{f_1, f_2\}$, with $f_1 : \mathcal{X}_\lambda \rightarrow \{0, 1\}$ and $f_2(x) := \overline{f_1(x)}$. The adversary \mathcal{A} executes the setup algorithm $\text{Setup}(1^\lambda)$ to receive (mpk, msk) and generates a functional key $\text{sk}_{f_1} \leftarrow \text{KeyGen}(\text{msk}, f_1)$. It chooses a message $x \leftarrow \mathcal{X}_\lambda$ and sends $(\text{mpk}_1, \text{mpk}_2, \text{sk}_{f_1}, f_2, x, x)$ to the challenger, which uses mpk and x to compute $\text{ct} \leftarrow \text{Enc}(\text{mpk}, x)$. In the next step, the challenger computes $\text{Dec}(\text{sk}_{f_1}, f_2, \text{ct}) = f_1(x)$ (note that the scheme by design does not aim at verifying sk_{f_1} vs. f_2) but then then verification tests whether $f_1(x) = f_2(x)$. This check is always false, due to the definition of f_2 ($f_2(x) = \overline{f_1(x)} \neq f_2(x), \forall x \in \mathcal{X}_\lambda$), and gives us the consistency attack. \square

The CFE case. The analysis presented in the last section can be adapted to the case of CFE security. More precisely, we show that CFE security does not imply consistency, by presenting a scheme that is CFE secure but not consistent. The scheme is presented in Fig. 12, it is the CFE secure version of the brute force scheme as introduced in [23, Section 5] and further analyzed in [53].

Theorem 4.5. *Let $\text{PKE} = (\text{PKE.Setup}, \text{PKE.Enc}, \text{PKE.Dec})$ be an IND-CPA secure public-key encryption scheme, then the functional encryption scheme $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ in Fig. 12 is CFE secure in the random oracle model for H for any functionality class \mathcal{F} of polynomial size s (in the security parameter). Namely, for any PPT adversary \mathcal{A} there is a PPT simulator \mathcal{S} such that*

$$\text{Adv}_{\text{FE}, \mathcal{A}, \mathcal{S}}^{\text{CFE}}(\lambda) \leq \text{negl}(\lambda) .$$

Proof. We refer to [53] for a security proof of the construction. \square

We show that this scheme does not imply input consistency.

Theorem 4.6. *Let $\text{PKE} = (\text{PKE.Setup}, \text{PKE.Enc}, \text{PKE.Dec})$ be an IND-CPA secure public-key encryption scheme, then the functional encryption scheme $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ in Fig. 12 is not input consistent for a concrete functionality class of size $s = 2$. Namely, for any algorithm Extract , there exists a PPT adversary \mathcal{A} such that*

$$\Pr[\text{in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract}) = 1] = 1 .$$

Proof. The attack described in the proof of Theorem 4.3 also applies here, since, for sk_{f_i} , the scheme in Fig. 12 still performs a simple decryption at position i and hence will produce inconsistent outputs. \square

The scheme described in Fig. 12 is also not setup consistent.

Theorem 4.7. *Let $\text{PKE} = (\text{PKE.Setup}, \text{PKE.Enc}, \text{PKE.Dec})$ be an IND-CPA secure public-key encryption scheme, then the functional encryption scheme $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ in Fig. 12 is not setup consistent for a concrete functionality class of size $s = 2$. Namely, there exists a PPT adversary \mathcal{A} such that*

$$\Pr[\text{set-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}) = 1] = 1 .$$

Proof. The attack described in the proof of Theorem 4.3 still applies here, since the scheme in Fig. 12 does not verify the claim on the function to be decrypted (and simply takes the matching dimension). \square

<p><u>Setup(1^λ) :</u> For $i = 1, \dots, s$ run $(\text{pk}_i, \text{sk}_i) \leftarrow \text{PKE.Setup}(1^\lambda)$ $(\text{mpk}, \text{msk}) = (\{\text{pk}_i\}_{i \in [s]}, \{\text{sk}_i\}_{i \in [s]})$ <u>KeyGen(msk, f_i) :</u> Return $\text{sk}_{f_i} = \text{sk}_i$ <u>Enc(mpk, x) :</u> Parse $\text{mpk} = \{\text{pk}_i\}_{i \in [n]}$ Sample $r_i \leftarrow \mathcal{Y}_\lambda$, for all $i \in [s]$ Compute $\text{ct}_i = (\text{Enc}(\text{mpk}, r_i), \text{H}(r_i) \oplus f_i(x)), \forall i \in [s]$ Return $\text{ct} = (\text{ct}_i)_{i \in [s]}$ <u>Dec($\text{sk}_{f_i}, f_i, (\text{ct}_i)_{i \in [s]}$) :</u> Parse $\text{ct}_i = (\text{ct}_{i,1}, \text{ct}_{i,2})$ Compute $r_i = \text{PKE.Dec}(\text{sk}_{f_i}, \text{ct}_{i,1})$ Return $y := \text{H}(r_i) \oplus \text{ct}_{i,2}$</p>

Fig. 12: A CFE secure but inconsistent functional encryption scheme.

4.4 Consistency does not amplify Confidentiality

To conclude the relationship graph in Fig. 8, we analyze if consistency allows to lift security, i.e., whether consistency coupled with IND-CPA would directly yield (any of the) IND-CCA or CFE security notions. Both of these results are answered in the negative in this section by showing that in general, malleability and consistency are not contradicting requirements as can be seen by existing ordinary PKE schemes (cast as special cases of FE).

We provide an explicit proof of this insight for strong input consistency. For concreteness, let R be an (efficiently computable) map on the plaintext space and let $\text{maul}_{\text{FE}}^R$ be (an efficiently computable) map such that for all plaintexts x and public parameters mpk in the range of Setup , and for any fixed randomness r , $\text{maul}_{\text{FE}}^R(\text{Enc}(\text{mpk}, x; r), \text{mpk}) = \text{Enc}(\text{mpk}, R(x); r')$ for some randomness string r' . We call the map R separating for a function $f \in \mathcal{F}$ (or f -separating for short), if the composed map $f \circ R : \mathcal{X} \rightarrow \mathcal{Y}$ is an injective map.

Theorem 4.8. *If a functional encryption scheme $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ for a functionality class \mathcal{F} admits an efficiently computable map $\text{maul}_{\text{FE}}^R$ for a plaintext map R that is f -separating (as defined above) for a given $f \in \mathcal{F}$ then the scheme cannot be CCA-secure. Furthermore, such CPA-secure schemes and concrete functionality classes exist in the standard model (under computational assumptions) which satisfy strong input consistency but are neither CFE-secure nor CCA-secure.*

Proof. To prove the first part we construct a generic attack given the assumptions on R : the adversary does never invoke its oracle KeyGen , picks two challenges $x_0 \neq x_1$ of the same length and obtains the challenge ciphertext ct^β as the encryption of m^β . The adversary can now query the decryption oracle for function f on $\text{maul}_{\text{FE}}^R(\text{ct}^\beta, \text{mpk})$ to obtain the function value y'^β (by the perfect correctness of the scheme). Since R is f -separating, $y'^\beta = f(R(x^\beta)) \neq f(R(x^{1-\beta}))$ and thus β can be guessed perfectly.

To prove the second part of the scheme we cast the El Gamal encryption scheme as an FE scheme with function class $\mathcal{F} = \{\text{id}, f_0\}$ which is therefore CPA-secure under DDH [23]: Let $G = \langle g \rangle$ be a prime-order group (for a prime $2^{\lambda-1} < q < 2^\lambda$) with generator g . More concretely, we let $(\text{mpk}, \text{msk}) \leftarrow (g^a, a)$ for a random exponent a ; an encryption of x is defined as $(g^r, g^{r_a}x)$ for a random exponent r ; and finally, define $\text{Dec}(\text{mpk}, \text{sk}_{\text{id}} = \text{msk} = a, \text{ct}' = (\text{ct}'_0, \text{ct}'_1))$ to return \diamond if $g^a \neq \text{mpk}$, and otherwise to return $x' \leftarrow \text{ct}'_1 \cdot (\text{ct}'_0)^{-a}$. The scheme satisfies strong input consistency, since given the ciphertext and the public-private key-pair

<p>Setup(1^λ) :</p> $\mathcal{G} := (\mathbb{G}, p, g) \leftarrow \text{GGen}(1^\lambda),$ $\mathbf{A} \leftarrow \mathcal{D}_k, \mathbf{W} \leftarrow \mathbb{Z}_p^{m \times (k+1)}$ $\text{mpk} := (\mathcal{G}, [\mathbf{A}]_g, [\mathbf{W}\mathbf{A}]_g), \text{msk} := \mathbf{W}$ Return (mpk, msk)	<p>KeyGen(mpk, msk, $\mathbf{y} \in \mathbb{Z}_Y^m, \mathbf{y} \in \mathbb{Z}_p^m$) :</p> Return $\text{sk}_{\mathbf{y}} := \mathbf{W}^\top \mathbf{y} \in \mathbb{Z}_p^{k+m+1}$
<p>Enc(mpk, $\mathbf{x} \in \mathbb{Z}_X^m, \mathbf{x} \in \mathbb{Z}_p^m$) :</p> $\mathbf{r} \leftarrow \mathbb{Z}_p^k, \begin{pmatrix} \mathbf{c}_0 \\ \mathbf{c}_1 \end{pmatrix} := \begin{pmatrix} -\mathbf{A}\mathbf{r} \\ \mathbf{x} + \mathbf{W}\mathbf{A}\mathbf{r} \end{pmatrix}$ Return $\text{ct} := [(\mathbf{c}_i^0)]_g \in \mathbb{G}^{k+m+1}$	<p>Dec(mpk, $\mathbf{y}, \text{sk}_{\mathbf{y}}, \text{ct}$) :</p> Parse $\text{ct} := [(\mathbf{c}_i^0)]_g$ $C := [\mathbf{c}_1^\top \mathbf{y} + \mathbf{c}_0^\top \text{sk}_{\mathbf{y}}]_g$ <div style="border: 1px solid black; padding: 2px; display: inline-block;">Return C</div> Return $\log(C)$

Fig. 13: FE for the class $\mathcal{F}_p^{m, X, Y}$, \mathcal{F}_p^m , $\mathcal{F}_{p, L}^m$ and the predicate class \mathcal{P}_p^m based on the \mathcal{D}_k -MDDH assumption.

(g^a, a) , the underlying message is committed to. Furthermore, the scheme is malleable and the mapping $R : x \mapsto c \cdot x$ for a constant c is an injective mapping which is separating the identity function $\text{id} \in \mathcal{F}$. We conclude the proof of the second part of the theorem by observing that CFE security for this scheme is impossible by the impossibility result given in [53, Theorem 5.1]. \square

5 Consistency Analysis of Selected Functional Encryption Schemes

In this section, we analyze the single-input functional encryption schemes for the inner product functionality based on the MDDH assumption regarding input consistency. These schemes have been initially introduced for the DDH assumption in [8] and extended to the MDDH assumption in [5]. This analysis contains of two parts: The analysis for the bounded-norm functionality class $\mathcal{F}_p^{m, X, Y}$ and the functionality class \mathcal{F}_p^m .

We observe that some schemes for the inner product functionality seem to be input consistent, but without specific modifications they are not. Therefore, we analyze these schemes and the corresponding modifications for input consistency in this section. For both of the mentioned functionality classes we obtain negative results, i.e. the analyzed scheme is neither input consistent for the functionality classes of bounded-norm inner products nor for the inner products calculated over \mathbb{Z}_p^m . To prove this, we present an attack for both cases.

Beside this, we introduce a natural modification of the above functionality class and denote it by $\mathcal{F}_{P, L}^m$. It turns out, that when instantiating the MDDH scheme of Fig. 13 with the modified inner product functionality class that the resulting scheme is an input consistent functional encryption scheme. This is formally defined and proven in Section 5.2.

In addition to this, we present a modification of the inner-product scheme described in Fig. 13, which covers a more restricted functionality class, \mathcal{P}_p^m . This scheme again, is input consistent.

The standard tools and assumptions for this section can be found in Section 2.3.

5.1 Inconsistency of the Plain Schemes

We state the following theorem regarding the input consistency of the inner-product functional encryption scheme described in Fig. 13.

Theorem 5.1. *The functional encryption scheme FE described in Fig. 13 for the functionality class $\mathcal{F}_p^{m, X, Y}$ and \mathcal{F}_p^m , with p prime, is not input consistent. Namely, for any algorithm Extract, there exists a PPT adversary \mathcal{A} such that*

$$\Pr[\text{in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract}) = 1] = 1 .$$

We prove Theorem 5.1 by separating it into two Lemmas. First, in Lemma 5.2, we show that the scheme described in Fig. 13 is input inconsistent for the functionality class $\mathcal{F}^{m,X,Y}$. Second, in Lemma 5.3, we show that the same scheme is input inconsistent for the functionality class \mathcal{F}_p^m .

Lemma 5.2. *The functional encryption scheme FE for the functionality class $\mathcal{F}^{m,X,Y}$ described in Fig. 13 is not input consistent. Namely, for any algorithm Extract, there exists a PPT adversary \mathcal{A} such that:*

$$\Pr[\text{in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract}) = 1] = 1 .$$

Proof. For the computation of the final output in the decryption procedure, it is necessary to compute the discrete logarithm of $[\langle \mathbf{x}, \mathbf{y} \rangle]_g$. As described in [8], we assume that the computed inner product lies in an polynomial bounded interval $\{0, \dots, L\}$, i.e. $\langle \mathbf{x}, \mathbf{y} \rangle \in \{0, \dots, L\}$ with a known L . This ensures that the discrete logarithm computation can be performed in $\tilde{O}(L^{1/2})$, using Pollard's kangaroo method [58] (or even $\tilde{O}(L^{1/3})$, by precomputing a table of size $\tilde{O}(L^{1/3})$ [20]). Due to correctness, we assume that for every encrypted vector \mathbf{x} , with $\|\mathbf{x}\|_\infty < X$, and every functional key corresponding to \mathbf{y} , with $\|\mathbf{y}\|_\infty < Y$, the decryption gives us the right output $\langle \mathbf{x}, \mathbf{y} \rangle$. This results in the fact that L must be bigger than $m \cdot X \cdot Y$. In this case, the decryption procedure outputs \perp .

Now, we describe the behavior of an attacker \mathcal{A} against the input consistency of the scheme. After the challenger has generated the parameters, $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and has sent mpk to the adversary, the adversary generates a ciphertext ct , by encrypting the vector $\mathbf{x} := (L+1) \cdot \mathbf{e}_1$ after the rules defined in the encryption procedure. In the next step, \mathcal{A} queries the key generation oracle for the vectors \mathbf{e}_1 and \mathbf{e}_m , receives $\text{sk}_{\mathbf{e}_1}$ and $\text{sk}_{\mathbf{e}_m}$ as a reply and sends ct_x to the challenger.

After receiving ct , the challenger executes $\text{Extract}(\text{ct}, \{(\text{sk}_{\mathbf{e}_1}, \mathbf{e}_1), (\text{sk}_{\mathbf{e}_m}, \mathbf{e}_m)\})$ and computes $y_1 = \text{Dec}(\text{mpk}, \mathbf{e}_1, \text{sk}_{\mathbf{e}_1}, \text{ct})$ and $y_2 = \text{Dec}(\text{mpk}, \mathbf{e}_m, \text{sk}_{\mathbf{e}_m}, \text{ct})$. We consider the computation of y_1 and y_2 in more detail. In the decryption $\text{Dec}(\text{mpk}, \mathbf{e}_1, \text{sk}_{\mathbf{e}_1}, \text{ct})$, the decryptor computes $g^{(L+1) \cdot \mathbf{e}_1, \mathbf{e}_1} = g^{L+1}$ and tries to perform the discrete logarithm computation. This computation fails, due to the fact that $L+1$ is not part of the bounded interval $\{0, \dots, L\}$, therefore the procedure outputs \perp (this argument can be made for any fixed bound L). For the decryption procedure $\text{Dec}(\text{mpk}, \mathbf{e}_m, \text{sk}_{\mathbf{e}_m}, \text{ct})$, the decryptor computes $g^{(L+1) \cdot \mathbf{e}_1, \mathbf{e}_m} = g^0 = 1$, for which the discrete logarithm can be easily computed. This results in $y_1 := \perp$ and $y_2 = 0$.

For the consistency check, we need to compute the preimages of the two different encryptions, i.e. $f_{\mathbf{e}_1}^{-1}(\perp)$ and $f_{\mathbf{e}_m}^{-1}(0)$. The first preimage is defined as $f_{\mathbf{e}_1}^{-1}(\perp) = \{\perp\}$ (due to Definition 2.2) and the second preimage as $f_{\mathbf{e}_m}^{-1}(0) = \{\mathbf{x} \in \mathbb{Z}_X^m : \langle \mathbf{x}, \mathbf{e}_m \rangle = 0\} = \{(\begin{smallmatrix} \mathbf{x} \\ 0 \end{smallmatrix}), \text{ with } \mathbf{x} \in \mathbb{Z}_X^{m-1}\}$. For the final step in the consistency check, we compute the intersection of the two preimages $f_{\mathbf{e}_1}^{-1}(\perp) \cap f_{\mathbf{e}_m}^{-1}(0) = \{\perp\} \cap \{(\begin{smallmatrix} \mathbf{x} \\ 0 \end{smallmatrix}), \text{ with } \mathbf{x} \in \mathbb{Z}_X^{m-1}\} = \emptyset$.

This results in a consistency attack regardless of the output of $\text{Extract}(\text{ct}, \{(\text{sk}_{\mathbf{e}_1}, \mathbf{e}_1), (\text{sk}_{\mathbf{e}_m}, \mathbf{e}_m)\})$. This shows, that the described adversary \mathcal{A} always wins the game for every algorithm Extract. \square

Lemma 5.3. *Let FE be the IND-CPA secure functional encryption scheme for the functionality class \mathcal{F}_p^m , with p prime, described in Fig. 13, then the scheme FE is not input consistent. Namely, for any algorithm Extract, there exists a PPT adversary \mathcal{A} such that:*

$$\Pr[\text{in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract}) = 1] = 1 .$$

Proof. The proof works in the same manner as for Lemma 5.2. The polynomial bound L for the discrete logarithm computation in the last step must be smaller than p , such that we can still find a value $L+1$ for which the described attack works. If this is not the case, and the decryption procedure still remains efficient, it is possible to compute the discrete logarithm of g^x for all $x \in \mathbb{Z}_p$ by letting the decryption algorithm perform the task on random group elements. This yields a contradiction against the MDDH assumption and therefore, due to the fact that the security of the scheme is based on MDDH, a contradiction against the IND-CPA security of the scheme. \square

5.2 Consistency for Inner-product Schemes

Now, we present the modified inner product functionality class and state the theorem that when instantiating the scheme in Fig. 13 for this functionality class it achieves input consistency. The main idea of the new

functionality class is that we allow the decryption procedure to output a new error symbol `oob` in the case that it is not able to do the discrete logarithm computation in the last step. The preimage of the `oob` symbol is then defined as all the \mathbf{x} such that $\langle \mathbf{x}, \mathbf{y} \rangle$ exceeds the polynomial bound necessary for the logarithm computation. This allows to prevent the input consistency attack described in the proof of Theorem 5.1.

Modified Inner Product over \mathbb{Z}_P . Let $\mathcal{F} = \{\mathcal{F}_{P_\lambda, L_\lambda}^m\}_{\lambda \in \mathbb{N}}$ be a family (indexed by λ) of sets $\mathcal{F}_{P_\lambda, L_\lambda}^m$, where P_λ is a modulus of length λ and L_λ . Omitting the index λ , the set $\mathcal{F}_{P, L}^m = \{f_{\mathbf{y}} : \mathbb{Z}_P^m \rightarrow \mathbb{Z}_P, \text{ for } \mathbf{y} \in \mathbb{Z}_P^m\}$ where

$$f_{\mathbf{y}}(\mathbf{x}) = \begin{cases} \langle \mathbf{x}, \mathbf{y} \rangle \bmod P & \text{if } \langle \mathbf{x}, \mathbf{y} \rangle \in \{0, \dots, L\} \\ \text{oob} & \text{if } \langle \mathbf{x}, \mathbf{y} \rangle > L. \end{cases}$$

defines the inner-product operation over \mathbb{Z}_P . In addition to that, we define an out-of-bound symbol `oob` that is defined as the output of the function when the resulting inner product computation does not lie within a polynomial bound $\{0, \dots, L\}$. The preimage of the `oob` symbol is defined as follows

$$f_{\mathbf{y}}^{-1}(\text{oob}) = \{\mathbf{x} \in \mathbb{Z}_P^m : \langle \mathbf{x}, \mathbf{y} \rangle > L\}.$$

The preimage for all other outcomes is defined as

$$f_{\mathbf{y}}^{-1}(z) = \{\mathbf{x} \in \mathbb{Z}_P^m : \langle \mathbf{x}, \mathbf{y} \rangle = z\}.$$

When we consider the functional encryption scheme for the modified inner-product functionality $\mathcal{F}_{p, L}^m$, it achieves input consistency.

Theorem 5.4. *The functional encryption scheme FE described in Fig. 13 for the functionality class $\mathcal{F}_{p, L}^m$, with p prime, is input consistent. Namely, for any PPT adversary \mathcal{A} , it holds that*

$$\Pr[\text{in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract}) = 1] = 0,$$

where `Extract` is the algorithm that always returns unknown.

Proof. To prove the input consistency of the scheme described in Fig. 13, we need to show that no matter what ciphertext an adversary generates there exists at least one underlying plaintext that explains the decryption behavior of `ct` under different functional keys $\text{sk}_{\mathbf{y}_i}$ queried by the adversary \mathcal{A} during the game. We prove this by relying on the algebraic properties of the groups for which the functional encryption scheme is defined. In more detail, we show that there exists always a solution for a linear equation system that is defined by the different inner product computations between the functional keys and the submitted ciphertext. The existence of a solution shows that there exists at least one plaintext that explains the functional decryption behavior.

Now, we describe how the game proceeds. In the first step, the challenger generates the master public key and the master secret key by executing the setup procedure $(\text{mpk}, \text{msk}) = ((\mathcal{G}, [\mathbf{A}]_g, [\mathbf{W}\mathbf{A}]_g), \mathbf{W}) \leftarrow \text{Setup}(1^\lambda)$. In the next step, the adversary \mathcal{A} receives `mpk` and has access to a key generation oracle $\text{KeyGen}(\text{mpk}, \text{msk}, \cdot)$. Whenever \mathcal{A} queries the key generation oracle with a vector \mathbf{y}_i , the challenger generates $\text{sk}_{\mathbf{y}_i} = \text{KeyGen}(\text{mpk}, \text{msk}, \mathbf{y}_i)$, adds $(\text{sk}_{\mathbf{y}_i}, \mathbf{y}_i)$ to the list F and sends $\text{sk}_{\mathbf{y}_i}$ to \mathcal{A} . At some point in the game, \mathcal{A} sends `ct` to the challenger and the challenger computes $[z_i]_g := \text{Dec}(\text{mpk}, \mathbf{y}_i, \text{sk}_{\mathbf{y}_i}, \text{ct})$ for all $(\text{sk}_{\mathbf{y}_i}, \mathbf{y}_i) \in F$. We consider how the decryption works in more detail and determine $[z_i]_g$ specifically corresponding to \mathbf{y}_i .

For the ciphertext, output by the adversary \mathcal{A} , we write $\text{ct} = \begin{pmatrix} \mathbf{c}'_0 \\ \mathbf{c}'_1 \end{pmatrix}$, with $\mathbf{c}'_0 = [\mathbf{c}_0]_g \in \mathbb{G}^{k+1}$ and $\mathbf{c}'_1 = [\mathbf{c}_1]_g \in \mathbb{G}^m$. To be more specific, we also write $\mathbf{c}'_0 \in \mathbb{G}^{k+1}$ and $\mathbf{c}'_1 \in \mathbb{G}^m$ as explicit group elements, i.e. $\mathbf{c}'_0 := \begin{pmatrix} g^{c_{0,1}} \\ \vdots \\ g^{c_{0,k+1}} \end{pmatrix}$ and $\mathbf{c}'_1 := \begin{pmatrix} g^{c_{1,1}} \\ \vdots \\ g^{c_{1,m}} \end{pmatrix}$ with the generator g , $\mathbf{c}_0 := (c_{0,1}, \dots, c_{0,k+1}) \in \mathbb{Z}_p^{k+1}$ and $\mathbf{c}_1 := (c_{1,1}, \dots, c_{1,m}) \in \mathbb{Z}_p^m$. To show that the decryption procedure always decrypts to one underlying element, we compute the decryption procedure for an arbitrary honestly generated key $\text{sk}_{\mathbf{y}}$. We denote the master secret key as $\mathbf{W} := \begin{pmatrix} | & & | \\ \mathbf{w}_1 & \dots & \mathbf{w}_{k+1} \\ | & & | \end{pmatrix}$, and correspondingly $\mathbf{W}^\top := \begin{pmatrix} - & \mathbf{w}_1^\top & - \\ & \vdots & \\ - & \mathbf{w}_{k+1}^\top & - \end{pmatrix}$. Using the matrix description,

the functional key is defined as $\mathbf{W}^\top \cdot \mathbf{y} = \begin{pmatrix} \langle \mathbf{w}_1^\top, \mathbf{y} \rangle \\ \vdots \\ \langle \mathbf{w}_{k+1}^\top, \mathbf{y} \rangle \end{pmatrix} \in \mathbb{Z}_p^{k+1}$. For the decryption, we need to compute two

different components: $[\mathbf{c}_0^\top \mathbf{sk}_\mathbf{y}]_g$ and $[\mathbf{c}_1^\top \mathbf{y}]_g$.

First, we describe how to compute $[\mathbf{c}_0^\top \mathbf{sk}_\mathbf{y}]_g$: We exponentiate all of the components of \mathbf{c}'_1 with the components of $\mathbf{sk}_\mathbf{y}$ and compute the product of the resulting vector components. In more detail, we compute $\prod_{i \in [k+1]} g^{c_{0,i} \cdot \langle \mathbf{w}_i^\top, \mathbf{y} \rangle} = g^{\sum_{i \in [k+1]} c_{0,i} \cdot \langle \mathbf{w}_i^\top, \mathbf{y} \rangle} = g^{\langle \mathbf{W} \cdot \mathbf{c}_0, \mathbf{y} \rangle}$. We proceed in the same way for the second component $\prod_{i \in [m]} g^{c_{1,i} \cdot y_i} = g^{\sum_{i \in [m]} c_{1,i} \cdot y_i} = g^{\langle \mathbf{c}_1, \mathbf{y} \rangle}$.

For the final computation, before the discrete logarithm computation, we need to multiply the two components, which results in $g^{\langle \mathbf{W} \cdot \mathbf{c}_0, \mathbf{y} \rangle} \cdot g^{\langle \mathbf{c}_1, \mathbf{y} \rangle} = g^{\langle \mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1, \mathbf{y} \rangle} = g^{\langle \mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1, \mathbf{y} \rangle}$.

In the final step of the decryption procedure the discrete logarithm computation happens. We denote the final decryptions with respect to the different \mathbf{y}_i 's by $z_i := \log(g^{\langle \mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1, \mathbf{y} \rangle})$.

To prove the input consistency, we show that the preimage of z_i contains the value $\mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1$ for the case that $z_i = \text{oob}$ and the case that $z_i \neq \text{oob}$. This leads to the fact that $\mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1 \in \bigcap_{i \in [n]} f_{\mathbf{y}_i}^{-1}(z_i)$ for all $z_i \in \{0, \dots, L\} \cup \{\text{oob}\}$, which covers all the possible values of z_i .

Both of these cases follow directly from the definition of the preimage. In more detail, as described in the beginning of Section 5.2, it holds that $f_{\mathbf{y}}^{-1}(\text{oob})$ contains all the vectors \mathbf{x} such that $\langle \mathbf{x}, \mathbf{y} \rangle > L$. For the case that $z_i = \text{oob}$ it holds that $\langle \mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1, \mathbf{y}_i \rangle > L$, after the analysis above, and therefore $\mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1 \in f_{\mathbf{y}_i}^{-1}(\text{oob})$. For the case that $z_i \in \{0, \dots, L\}$ it holds that the preimage contains all the vectors \mathbf{x} , such that $\langle \mathbf{x}, \mathbf{y}_i \rangle = z_i$. Therefore, again with the analysis above, it follows that $\mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1 \in f_{\mathbf{y}_i}^{-1}(z_i)$ for $z_i \in \{0, \dots, L\}$. Overall, this leads to the fact that $\mathbf{W} \cdot \mathbf{c}_0 \in \bigcap_{i \in [n]} f_{\mathbf{y}_i}^{-1}(z_i)$ for all $i \in [n]$ with $z_i \in \{0, \dots, L\} \cup \{\text{oob}\}$. \square

The scheme described in Fig. 13 is obviously CPA secure for the functionality class $\mathcal{F}_{p,L}^m$ if the base FE scheme is CPA secure.

Theorem 5.5. *Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be the IND-CPA secure functional encryption scheme for the functionality class \mathcal{F}_p^m , with p prime, described in Fig. 13. Then the functional encryption scheme $\text{FE}' = (\text{Setup}', \text{KeyGen}', \text{Enc}', \text{Dec}')$ for the functionality class $\mathcal{F}_{p,L}^m$, with p prime, described in Fig. 13 is IND-CPA secure. Namely, for any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B} such that*

$$\text{Adv}_{\text{FE}', \mathcal{A}}^{\text{IND-CPA}}(\lambda) \leq \text{Adv}_{\text{FE}, \mathcal{B}}^{\text{IND-CPA}}(\lambda) .$$

5.3 Consistency of a related Predicate Scheme

The scheme used in this section, is another modified version of the inner product encryption scheme (described in Fig. 13) for the functionality class of inner product predicates \mathcal{P}_p^m . We define the functionality class more formally:

Inner Product Predicate over \mathbb{Z}_p . Let $\mathcal{P} = \{\mathcal{P}_{p,\lambda}^m\}_{\lambda \in \mathbb{N}}$ be a family of sets $\mathcal{P}_{p,\lambda}^m$, where p is a prime of length λ and \mathbb{G} a group of size p and generator g .¹¹ The sets are defined by $\mathcal{P}_p^m = \{f_{g,\mathbf{y}} : (\mathbb{Z}_p^m) \rightarrow \mathbb{G}, \text{ with } \mathbf{y} \in \mathbb{Z}_p^m\}$ where

$$f_{g,\mathbf{y}}(\mathbf{x}) = g^{\langle \mathbf{x}, \mathbf{y} \rangle} .$$

To modify the inner product encryption scheme to fit our new functionality class, we proceed without the discrete logarithm computation in the end of the decryption procedure and just output the value $g^{\langle \mathbf{x}, \mathbf{y} \rangle}$. The resulting scheme can now be used for the computation of inner-product predicates, e.g. for an inner product of 0 the decryption procedure would output 1. For this scheme, the input consistency property can be proven formally.

¹¹ As already mentioned in Section 2, we omit λ when it is clear from the context.

Theorem 5.6. *The functional encryption scheme FE for the functionality class \mathcal{P}_p^m , with p prime, described in Fig. 13 is input consistent. Namely, for any PPT adversary \mathcal{A} , it holds that*

$$\Pr[\text{in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract}) = 1] = 0 ,$$

where Extract is the algorithm that always returns unknown.

Proof. We proceed in a similar way as in the proof of Theorem 5.4. To prove the input consistency of the described scheme, we need to show that no matter what ciphertext an adversary generates there exists at least one underlying plaintext that would explain the decryption behavior of ct under different functional keys $\text{sk}_{\mathbf{y}_i}$ queried by the adversary \mathcal{A} during the game. We prove this by relying on the algebraic properties of the groups for which the functional encryption scheme is defined. In more detail, we show that there exists always a solution for a linear equation system that is defined by the different inner product computations between the functional keys and the submitted ciphertext. The existence of a solution shows that there exists at least one plaintext that explains the functional decryption behavior.

Now, we describe how the game proceeds. In the first step, the challenger generates the master public key and the master secret key by executing the setup procedure $(\text{mpk}, \text{msk}) = ((\mathcal{G}, [\mathbf{A}]_g, [\mathbf{WA}]_g), \mathbf{W}) \leftarrow \text{Setup}(1^\lambda)$. In the next step, the adversary \mathcal{A} receives mpk and has access to a key generation oracle $\text{KeyGen}(\text{mpk}, \text{msk}, \cdot)$. Whenever \mathcal{A} queries the key generation oracle with a vector \mathbf{y}_i , the challenger generates $\text{sk}_{\mathbf{y}_i} = \text{KeyGen}(\text{mpk}, \text{msk}, \mathbf{y}_i)$, adds $(\text{sk}_{\mathbf{y}_i}, \mathbf{y}_i)$ to the list F and sends $\text{sk}_{\mathbf{y}_i}$ to \mathcal{A} . At some point in the game, \mathcal{A} sends ct to the challenger and the challenger computes $[z_i]_g := \text{Dec}(\text{mpk}, \mathbf{y}_i, \text{sk}_{\mathbf{y}_i}, \text{ct})$ for all $(\text{sk}_{\mathbf{y}_i}, \mathbf{y}_i) \in F$. We consider how the decryption works in more detail and determine $[z_i]_g$ specifically corresponding to \mathbf{y}_i .

For the ciphertext, output by the adversary \mathcal{A} , we write $\text{ct} = \begin{pmatrix} c'_0 \\ c'_1 \end{pmatrix}$, with $\mathbf{c}'_0 = [c_0]_g \in \mathbb{G}^{k+1}$ and $\mathbf{c}'_1 = [c_1]_g \in \mathbb{G}^m$. To be more specific, we also write $\mathbf{c}'_0 \in \mathbb{G}^{k+1}$ and $\mathbf{c}'_1 \in \mathbb{G}^m$ as explicit group elements, i.e. $\mathbf{c}'_0 := \begin{pmatrix} g^{c_{0,1}} \\ \vdots \\ g^{c_{0,k+1}} \end{pmatrix}$ and $\mathbf{c}'_1 := \begin{pmatrix} g^{c_{1,1}} \\ \vdots \\ g^{c_{1,m}} \end{pmatrix}$ with the generator g , $\mathbf{c}_0 := (c_{0,1}, \dots, c_{0,k+1}) \in \mathbb{Z}_p^{k+1}$ and $\mathbf{c}_1 := (c_{1,1}, \dots, c_{1,m}) \in \mathbb{Z}_p^m$. To show that the decryption procedure always decrypts to one underlying element, we compute the decryption procedure for an arbitrary honestly generated key $\text{sk}_{\mathbf{y}}$. We denote the master secret key as $\mathbf{W} := \begin{pmatrix} | & & | \\ \mathbf{w}_1 & \dots & \mathbf{w}_{k+1} \\ | & & | \end{pmatrix}$, and correspondingly $\mathbf{W}^\top := \begin{pmatrix} - & \mathbf{w}_1^\top & - \\ \vdots & \vdots & \vdots \\ - & \mathbf{w}_{k+1}^\top & - \end{pmatrix}$. Using the matrix description,

the functional key is defined as $\mathbf{W}^\top \cdot \mathbf{y} = \begin{pmatrix} \langle \mathbf{w}_1^\top, \mathbf{y} \rangle \\ \vdots \\ \langle \mathbf{w}_{k+1}^\top, \mathbf{y} \rangle \end{pmatrix} \in \mathbb{Z}_p^{k+1}$. For the decryption, we need to compute two different components: $[c_0^\top \text{sk}_{\mathbf{y}}]_g$ and $[c_1^\top \mathbf{y}]_g$.

First, we describe how to compute $[c_0^\top \text{sk}_{\mathbf{y}}]_g$: We exponentiate all of the components of \mathbf{c}'_0 with the components of $\text{sk}_{\mathbf{y}}$ and compute the product of the resulting vector components. In more detail, we compute $\prod_{i \in [k+1]} g^{c_{0,i} \cdot \langle \mathbf{w}_i^\top, \mathbf{y} \rangle} = g^{\sum_{i \in [k+1]} \langle c_{0,i} \cdot \mathbf{w}_i^\top, \mathbf{y} \rangle} = g^{\langle \sum_{i \in [k+1]} c_{0,i} \cdot \mathbf{w}_i^\top, \mathbf{y} \rangle} = g^{\langle \mathbf{W} \cdot \mathbf{c}_0, \mathbf{y} \rangle}$. We proceed in the same way for the second component $\prod_{i \in [m]} g^{c_{1,i} \cdot y_i} = g^{\sum_{i \in [m]} c_{1,i} \cdot y_i} = g^{\langle \mathbf{c}_1, \mathbf{y} \rangle}$.

To generate the final output, we need to multiply the two components, which results in $g^{\langle \mathbf{W} \cdot \mathbf{c}_0, \mathbf{y} \rangle} \cdot g^{\langle \mathbf{c}_1, \mathbf{y} \rangle} = g^{\langle \mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1, \mathbf{y} \rangle} = g^{\langle \mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1, \mathbf{y} \rangle}$.

Coming back to the initial description of the game, the decryption procedure outputs $[z_i]_g := [\langle \mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1, \mathbf{y}_i \rangle]_g$. Due to the fact that the vectors \mathbf{c}_0 and \mathbf{c}_1 are set by the adversary and the matrix \mathbf{W} is fixed after the setup procedure, the decryption relies only on the value \mathbf{y}_i . This results in the decryptions $g^{\langle \mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1, \mathbf{y}_1 \rangle}, \dots, g^{\langle \mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1, \mathbf{y}_n \rangle}$ for all the different secret key queries \mathbf{y}_i . Consequently,

$$\mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1 \in f_{g, \mathbf{y}_1}^{-1}([z_1]_g), \dots, \mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1 \in f_{g, \mathbf{y}_n}^{-1}([z_n]_g), \text{ which further implies that}$$

$$\mathbf{W} \cdot \mathbf{c}_0 + \mathbf{c}_1 \in \bigcap_{i \in [n]} f_{g, \mathbf{y}_i}^{-1}([z_i]_g) .$$

This makes the intersection non-empty for every possible ciphertext ct generated by \mathcal{A} . Therefore, the scheme is input consistent. \square

The scheme described in Fig. 13 for the functionality class \mathcal{P}_p^m achieves IND-CPA security:

Theorem 5.7. *Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be the IND-CPA secure functional encryption scheme for the functionality class \mathcal{F}_p^m , with p prime, described in Fig. 13. Then the functional encryption scheme*

$\text{FE}' = (\text{Setup}', \text{KeyGen}', \text{Enc}', \text{Dec}')$ for the functionality class \mathcal{P}_p^m , with p prime, described in Fig. 13 is IND-CPA secure. Namely, for any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B} such that

$$\text{Adv}_{\text{FE}', \mathcal{A}}^{\text{IND-CPA}}(\lambda) \leq \text{Adv}_{\text{FE}, \mathcal{B}}^{\text{IND-CPA}}(\lambda) .$$

This statement follows by a straightforward reduction to CPA security by observing that the restriction on the functional keys, i.e., the requirement $f_{\mathbf{y}}(\mathbf{x}^0) = f_{\mathbf{y}}(\mathbf{x}^1)$ is preserved for all keys, because if $\langle \mathbf{x}^0, \mathbf{y} \rangle = \langle \mathbf{x}^1, \mathbf{y} \rangle$ then it follows that $g^{\langle \mathbf{x}^0, \mathbf{y} \rangle} = g^{\langle \mathbf{x}^1, \mathbf{y} \rangle}$.

6 Consistency Compilers

In this section, we present black-box compilers that achieve consistency under the different corruption sets. The goal is to apply the compiler on top of a functional encryption scheme to achieve the desired consistency notion. For the case of CPA and CFE security, we present compilers that preserve the security of the underlying functional encryption scheme. For the case of CCA security, we present so called “advanced” compilers that lift the security of a CPA secure scheme into a CCA secure and consistent scheme. We note in passing that establishing the Naor-Yung security lifting in the FE space is of independent interest. We provide the first rigorous proof that the technique generalizes to the FE setting with a couple of technical modifications compared to the PKE case. The section is organized as follows:

- In Section 6.1, we present a first compiler that takes a CPA or CFE secure FE scheme and makes it a CPA resp. CFE-secure input consistent FE scheme. Input consistency is proven by Theorem 6.1 and security preservation for CPA and CFE in Theorem 6.2 and Theorem 6.5, respectively. Next, we define an advanced compiler that lifts any CPA secure FE scheme to an input consistent CCA secure FE scheme. Input consistency is proven in Theorem 6.10 and security lifting in Theorem 6.6.
- In Section 6.2, we present one compiler that takes a CPA or CFE secure FE scheme and makes it setup consistent, Theorem 6.11, and preserves CPA security, Theorem 6.12, and CFE security under natural restrictions, Theorem 6.13. As above, we show that CCA secure setup-consistency is achievable by actually presenting a second compiler that takes any CPA secure FE scheme and compiles it into a setup consistent CCA secure FE scheme. This is proven in Theorem 6.14 and in Theorem 6.19.
- Finally, in Section 6.3 we present a construction that achieves strong input consistency based on any verifiable FE scheme in a black-box way as proven in Theorem 6.20 (preserving the confidentiality notion of the underlying scheme).

6.1 Input Consistency

First Compiler. To achieve input-consistency under CPA and CFE security, we augment the output of an encryption algorithm with a non-interactive zero-knowledge (NIZK) proof that an underlying plaintext exists. The NIZK proof is generated over the master public key, the encryption algorithm’s randomness and the underlying plaintext. The zero-knowledge property of the NIZK makes sure that no information about the underlying plaintext is leaked, whereas the soundness prevents a malicious party from generating a valid proof over an invalid ciphertext. A formal description of this compiler is presented in Fig. 14 and the relation R_{in} , that needs to be supported by the NIZK scheme, is defined in Fig. 15. We show that the described construction indeed turns a functional encryption scheme into an input consistent functional encryption scheme.

```

Setup'(1λ) :
CRS ← NIZK.Setup(1λ)
(mpk, msk) ← Setup(1λ)
Return (mpk', msk') = ((CRS, mpk), msk)
KeyGen'(mpk', msk', f) :
Parse mpk' = (CRS, mpk), msk' = msk
skf = KeyGen(mpk, msk, f)
Return sk'f = skf
Enc'(mpk', x) :
Parse mpk' = (CRS, mpk)
ct = Enc(mpk, x; r) with r ← {0, 1}λ
Generate π ← NIZK.Prove(CRS, (mpk, ct), (x, r)) for Rin (Fig. 15)
Return ct' = (ct, π)
Dec'(mpk', f, sk'f, ct') :
Parse mpk' := (CRS, mpk), sk'f = skf, ct' = (ct, π)
If NIZK.Verify(CRS, (mpk, ct), π) = 1
    Return Dec(mpk, f, skf, ct)
Else
    Return ⊥

```

Fig. 14: Input consistency compiler

```

Relation Rin :
Instance: z = (mpk, ct)
Witness: w = (x, r), x ∈ X, random coins r
Rin(z, w) = 1 if and only if:
    ct = Enc(mpk, x; r)

```

Fig. 15: Relation used in the input consistency compiler

Theorem 6.1. *Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be a functional encryption scheme and $\text{NIZK} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ a NIZK proof system for relation R_{in} , then the construction $\text{FE}' = (\text{Setup}', \text{KeyGen}', \text{Enc}', \text{Dec}')$ defined in Figure 14 satisfies input consistency. Namely, for any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B} such that:*

$$|\Pr[\text{in-CONS}^{\text{FE}'}(1^\lambda, \mathcal{A}, \text{Extract}) = 1]| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}}^{\text{Sound}}(\lambda),$$

where Extract is the algorithm that always returns unknown.

Proof. To prove the input consistency of the scheme FE' , we rely on the soundness of the NIZK proof system. In more detail, we construct an adversary \mathcal{B} that generates a malicious proof, by relying on an adversary \mathcal{A} for the input consistency experiment $\text{in-CONS}^{\text{FE}}$.

In the beginning of the reduction, \mathcal{B} receives a common reference string CRS from its underlying experiment, it generates $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$, sets $\text{mpk}' := (\text{CRS}, \text{mpk})$ and sends mpk' to \mathcal{A} .

Whenever \mathcal{A} asks a key generation query f , \mathcal{B} computes the key $\text{sk}_f \leftarrow \text{KeyGen}(\text{msk}, f)$, adds (sk_f, f) to the list F and sends sk_f to \mathcal{A} .

At some point, \mathcal{A} sends a ciphertext $\text{ct}' = (\text{ct}, \pi)$ to \mathcal{B} . If $\text{NIZK.Verify}((\text{mpk}, \text{ct}), \pi) = 0$ then the adversary \mathcal{B} halts. In this case, $\text{Dec}(\text{mpk}, f_i, \text{sk}_{f_i}, \text{ct})$ yields \perp for all $i \in [n]$, by definition of the compiler (and hence the adversary \mathcal{A} loses the game). If $\text{NIZK.Verify}((\text{mpk}, \text{ct}), \pi) = 1$, \mathcal{B} simply outputs $(\text{mpk}_1, \text{ct}, \pi)$ as its forgery and halts.

Let us analyze the output of \mathcal{A} to see that the condition to break input consistency must imply a soundness violation of the NIZK scheme. In more detail, we define the event E as the event that the adversary \mathcal{A} performs a consistency attack under the assumption that $(\text{mpk}, \text{ct}) \in L$ and show that the occurrence of the event E would contradict the assumption.

Now, we analyze the possible outcomes for the decryptions y_i in the case of a consistency attack. We show that $y_i \neq \perp$ for all $i \in [n]$ (this is covered by event E_1). Furthermore we show that if $y_i \neq \perp$ then there exists an x such that $x \in \bigcap_{i \in [n]} f_i^{-1}(y_i)$ (this is denoted by event E_2).

In the case of event E_1 , we assume that at least one of the decryptions is equal to \perp , i.e. $y_i \neq \perp$. We distinguish between two cases:

1. It holds that $y_i = \perp$ for all $i \in [n]$. In this case, the adversary \mathcal{A} did not perform a consistency attack. In more detail, the intersection $\bigcap_{i \in [n]} f_i^{-1}(y_i)$ will contain the \perp value.
2. At least one, but not all, of the decryptions are equal to \perp , i.e. $y_i = \perp$. Since \perp is not an element \mathcal{X} , and therefore not an encryption value, then, by perfect correctness of the underlying FE scheme, it follows that there exists no w such that $((\text{mpk}, \text{ct}), w) \in R_{\text{in}}$ (i.e. it cannot be a valid instance). This is a contradiction to the assumption that $(\text{mpk}, \text{ct}) \in L$.

Considering both the above mentioned points, we can conclude that $y_i \neq \perp$ for all $i \in [n]$.

For the analysis of event E_2 , we assume, for the sake of contradiction, that the intersection $\bigcap_{i \in [n]} f_i^{-1}(y_i)$ is empty and it holds (with respect to event E_1) that $y_i \neq \perp$ for all $i \in [n]$. In this case, the adversary \mathcal{A} has generated a valid proof π for an $x \notin L_{\text{in}}$. Again, by the perfect correctness of the FE scheme, the adversary \mathcal{B} broke the soundness of the NIZK scheme, because it has found a ciphertext ct and provided a proof to be a valid encryption while the functional outputs say that there is no such underlying plaintext.

By combining the events E_1 and E_2 , we have proven that event E cannot occur. To recap, whenever $(\text{mpk}, \text{ct}) \in L$, it is not possible for an adversary \mathcal{A} to perform a consistency attack. Hence the only way the adversary can break setup consistency is by breaking the soundness property of the NIZK scheme, i.e., providing the statement $(\text{mpk}, \text{ct}) \notin L$.

This yields the bound

$$|\Pr[\text{in-CONS}^{\text{FE}'}(1^\lambda, \mathcal{A}, \text{Extract}) = 1]| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}}^{\text{Sound}}(\lambda) .$$

and therefore we obtain the theorem. □

Besides proving that the compiler achieves input consistency, we also need to prove the security preservation under the two different notions of IND-CPA security and CFE security.

We first prove the security preservation of the compiler under CPA security and conclude with the preservation for CFE security.

Theorem 6.2. *Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be an IND-CPA secure functional encryption scheme and $\text{NIZK} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ a NIZK proof system, then the construction $\text{FE}' = (\text{Setup}', \text{KeyGen}', \text{Enc}', \text{Dec}')$, defined in Figure 14, is IND-CPA secure. Namely, for any PPT adversary \mathcal{A} , there exist PPT algorithms \mathcal{B} and \mathcal{B}' such that*

$$\text{Adv}_{\text{FE}', \mathcal{A}}^{\text{IND-CPA}}(\lambda) \leq 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{B}}^{\text{ZK}}(\lambda) + \text{Adv}_{\text{FE}, \mathcal{B}'}^{\text{IND-CPA}}(\lambda) .$$

Proof. To prove this statement, we use a hybrid argument with the games defined in Fig. 16. Note that G_0 corresponds to the game $\text{IND-CPA}_0^{\text{FE}}(1^\lambda, \mathcal{A})$ and G_3 to the game $\text{IND-CPA}_1^{\text{FE}}(1^\lambda, \mathcal{A})$. This results in:

$$\text{Adv}_{\text{FE}', \mathcal{A}}^{\text{IND-CPA}}(1^\lambda) = |\text{Win}_{\mathcal{A}}^{\text{G}_0}(1^\lambda) - \text{Win}_{\mathcal{A}}^{\text{G}_3}(1^\lambda)| .$$

We describe the different games in more detail:

Game G_1 : In this game, we change from an honestly generated CRS and honestly generated proofs to a simulated CRS and simulated proofs. The transition from G_0 to G_1 is justified by the zero-knowledge property of the NIZK. Namely, in Lemma 6.3, we exhibit a PPT adversary \mathcal{B}_0 such that:

$$|\text{Win}_{\mathcal{A}}^{\text{G}_0}(1^\lambda) - \text{Win}_{\mathcal{A}}^{\text{G}_1}(1^\lambda)| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_0}^{\text{ZK}}(\lambda) .$$

Game G_2 : In this game, we change from an encryption of x^0 to x^1 for the encryption queries. The transition from G_1 to G_2 is justified by the IND-CPA security of FE. Namely, in Lemma 6.4, we exhibit a PPT adversary \mathcal{B}_1 such that:

$$|\text{Win}_{\mathcal{A}}^{\text{G}_1}(1^\lambda) - \text{Win}_{\mathcal{A}}^{\text{G}_2}(1^\lambda)| \leq \text{Adv}_{\text{FE}, \mathcal{B}_1}^{\text{IND-CPA}}(\lambda) .$$

Game G_3 : This game is the $\text{IND-CPA}_1^{\text{FE}}(1^\lambda, \mathcal{A})$ game. The transition from G_2 to G_3 is almost symmetric to the transition from G_0 to G_1 except from the fact that the reduction encrypts x^1 instead of x^0 . As in Lemma 6.3, the transition is justified by the zero-knowledge property of NIZK. Namely, we can exhibit a PPT adversary \mathcal{B}_0 such that:

$$|\text{Win}_{\mathcal{A}}^{\text{G}_2}(1^\lambda) - \text{Win}_{\mathcal{A}}^{\text{G}_3}(1^\lambda)| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_0}^{\text{ZK}}(\lambda) .$$

Putting everything together, we obtain the theorem. \square

Lemma 6.3 (Transition from G_0 to G_1). *For any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B}_0 such that*

$$|\text{Win}_{\mathcal{A}}^{\text{G}_0}(1^\lambda) - \text{Win}_{\mathcal{A}}^{\text{G}_1}(1^\lambda)| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_0}^{\text{ZK}}(\lambda) .$$

Proof. We build an adversary \mathcal{B}_0 that simulates G_β towards \mathcal{A} when interacting with the underlying $\text{ZK}_\beta^{\text{NIZK}}$ experiment.

In the beginning of the reduction, \mathcal{B}_0 receives CRS from the $\text{ZK}_\beta^{\text{NIZK}}$ experiment. It generates a functional encryption instance $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$, sets $\text{mpk}' = (\text{CRS}, \text{mpk})$ and gives mpk' to the adversary.

Whenever \mathcal{A} asks an encryption query (x^0, x^1) , \mathcal{B}_0 generates the ciphertext $\text{ct} = \text{Enc}(\text{mpk}, x^0; r)$ with $r \leftarrow \{0, 1\}^\lambda$ and sends $y = (\text{mpk}, \text{ct})$ and $w = (x, r)$ as a statement-witness pair to its challenger. As an answer, \mathcal{B}_0 receives a proof π for R_{in} . It sets $\text{ct}' = (\text{ct}, \pi)$ and sends it to \mathcal{A} .

For a key generation query f , \mathcal{B}_0 generates $\text{sk}_f \leftarrow \text{KeyGen}(\text{mpk}, \text{msk}, f)$ for and sends $\text{sk}'_f = \text{sk}_f$ as a reply to \mathcal{A} .

This covers the simulation of the game G_β . Finally \mathcal{B}_0 outputs the same bit β' returned by \mathcal{A} . It follows, from the perfect simulation, that the advantage of \mathcal{B}_0 is the same as the advantage of \mathcal{A} . \square

Lemma 6.4 (Transition from G_1 to G_2). *For any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B}_1 such that*

$$|\text{Win}_{\mathcal{A}}^{G_1}(1^\lambda) - \text{Win}_{\mathcal{A}}^{G_2}(1^\lambda)| \leq \text{Adv}_{\text{FE}, \mathcal{B}_1}^{\text{IND-CPA}}(\lambda) .$$

Proof. We build an adversary \mathcal{B}_1 that simulates $G_{1+\beta}$ towards \mathcal{A} when interacting with the underlying $\text{IND-CPA}_{\beta}^{\text{FE}}$ experiment.

In the beginning of the reduction, \mathcal{B}_1 receives mpk from the experiment. It simulates a CRS, i.e. $(\text{CRS}, \tau) \leftarrow \mathcal{S}_1(1^\lambda)$, sets $\text{mpk}' = (\text{CRS}, \text{mpk})$ and gives mpk' to the adversary.

Whenever \mathcal{A} asks an encryption query (x^0, x^1) , \mathcal{B}_1 forwards it to its own encryption oracle to receive $\text{ct} \leftarrow \text{Enc}(\text{mpk}, x^\beta)$, simulates a proof for the relation R_{in} , i.e. $\pi \leftarrow \mathcal{S}_2(\text{CRS}, \tau, x^\beta)$ and sends $\text{ct}' = (\text{ct}, \pi)$ to \mathcal{A} .

For a key generation query f , \mathcal{B}_1 queries its own key generation oracle on f to receive $\text{sk}_f \leftarrow \text{KeyGen}(\text{mpk}, \text{msk}, f)$, sets $\text{sk}'_f = \text{sk}_f$ and sends sk'_f to \mathcal{A} .

This covers the simulation of the game $G_{1+\beta}$. Finally \mathcal{B}_1 outputs the same bit β' returned by \mathcal{A} . It follows, from the perfect simulation, that the advantage of \mathcal{B}_1 is the same as the advantage of \mathcal{A} . \square

Beside showing the IND-CPA security preservation, we also need to show the CFE security preservation.

Theorem 6.5. *Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be a CFE secure functional encryption scheme, i.e., there exists a simulator $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3)$ such that $\text{Real}^{\text{FE}}(1^\lambda, \mathcal{A}) \approx \text{Ideal}^{\text{FE}}(1^\lambda, \mathcal{A}, \mathcal{S})$ w.r.t. any adversary \mathcal{A} , and let further $\text{NIZK} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ be a NIZK proof system for the relation R_{in} , then for the construction $\text{FE}' = (\text{Setup}', \text{KeyGen}', \text{Enc}', \text{Dec}')$, defined in Figure 14, we can design a simulator $\mathcal{S}' = (\mathcal{S}'_1, \mathcal{S}'_2, \mathcal{S}'_3)$ such that for any adversary $\mathcal{A}' = (\mathcal{A}'_1, \mathcal{A}'_2)$ against the new scheme we can design adversaries \mathcal{A} and \mathcal{B} such that*

$$\text{Adv}_{\text{FE}', \mathcal{A}', \mathcal{S}'}^{\mathcal{D}', \text{CFE}}(\lambda) \leq \text{Adv}_{\text{FE}, \mathcal{A}}^{\mathcal{D}, \text{CFE}}(\lambda) + \text{Adv}_{\text{NIZK}, \mathcal{B}}^{\text{ZK}}(\lambda).$$

Proof. The simulator is defined as follows: \mathcal{S}'_1 runs \mathcal{S}_1 to obtain the pair (mpk, s) and runs the simulator of the NIZK scheme to obtain (CRS, τ) and outputs $\text{mpk}' := (\text{mpk}, \text{CRS})$. Next, $\mathcal{S}'_2 := \mathcal{S}_2$, and finally, \mathcal{S}'_3 runs \mathcal{S}_3 , receives a ciphertext ct_ℓ , and simulates a NIZK proof using the trapdoor τ for instance (mpk', ct) . Note that the internal state managed by \mathcal{S} is managed by \mathcal{S}' . We obtain the statement again by a sequence of hybrid steps. Let G_0 be the real CFE experiment and G_2 the ideal CFE experiment with the above simulator. Let G_1 be a hybrid experiment, where the only change is that we replace the CRS and simulate all NIZK proofs. Analogous to the proof of Theorem 6.2, an adversary \mathcal{B} with advantage α in distinguishing the outputs of experiments G_0 and G_1 (with respect to a certain adversary \mathcal{A}'') directly yields a distinguisher telling apart simulated and genuine proofs with the same advantage. For the second step, we see that any pair $(\mathcal{D}', \mathcal{A}')$ with advantage α in distinguishing the outputs of the experiments G_1 and G_2 can be transformed into a pair $(\mathcal{D}, \mathcal{A})$ such that the outputs of $\text{Real}^{\text{FE}}(1^\lambda, \mathcal{A})$ and $\text{Ideal}^{\text{FE}}(1^\lambda, \mathcal{A}, \mathcal{S})$ are distinguishable by \mathcal{D} with advantage α . To see this, note that by the modular design of the scheme, the adversary can be defined as follows: when the first adversary \mathcal{A}_1 receives mpk it simulates CRS and internally runs an instance of $\mathcal{A}'_1((\text{mpk}, \text{CRS}))$. For requests to the key-generation oracle, \mathcal{A}_1 simply relays them to the oracle of \mathcal{A}'_1 . Besides this, any internal state of \mathcal{A}'_1 is maintained by \mathcal{A}_1 , and passed on to the second adversary. \mathcal{A}_1 outputs whatever \mathcal{A}'_1 outputs. Second, whenever \mathcal{A}_2 receives some ciphertext, say ct_ℓ , it internally runs \mathcal{A}'_2 on input (ct_ℓ, π) , where π is a simulated proof for the relation R_{in} . Finally, \mathcal{A}_2 outputs whatever \mathcal{A}'_2 outputs. We see that the output distribution of $\text{Real}^{\text{FE}}(1^\lambda, \mathcal{A})$ is identical to the output of \mathcal{A}' in experiment G_1 and the output distribution of $\text{Ideal}^{\text{FE}}(1^\lambda, \mathcal{A}, \mathcal{S})$ is identical to the output distribution of G_2 , the ideal experiment with FE' and simulator \mathcal{S}' . This proves the theorem. \square

Second Advanced Compiler. For the advanced input consistency compiler that takes a CPA secure scheme and achieves CCA security, we make use of the Naor-Yung approach [55] and combine it with the approach of the presented input consistency compiler. In more detail, we run two different instances of the functional encryption scheme and create a proof that shows that both of these encryptions are generated in a

Game	CRS & π	ct	justification/remark
G_0	$\text{CRS} \leftarrow \text{NIZK.Setup}(1^\lambda)$ $\pi \leftarrow \text{NIZK.Prove}(\text{CRS}, x, w)$	$\text{Enc}(\text{mpk}, x^0)$	
G_1	$\text{CRS} \leftarrow \mathcal{S}_1(1^\lambda)$ $\pi \leftarrow \mathcal{S}_2(\text{CRS}, \tau, x)$	$\text{Enc}(\text{mpk}, x^0)$	Zero-knowledge of NIZK
G_2	$\text{CRS} \leftarrow \mathcal{S}_1(1^\lambda)$ $\pi \leftarrow \mathcal{S}_2(\text{CRS}, \tau, x)$	$\text{Enc}(\text{mpk}, \boxed{x^1})$	IND-CPA security of FE
G_3	$\text{CRS} \leftarrow \text{NIZK.Setup}(1^\lambda)$ $\pi \leftarrow \text{NIZK.Prove}(\text{CRS}, x, w)$	$\text{Enc}(\text{mpk}, x^1)$	Zero-knowledge of NIZK

Fig. 16: Overview of the games to prove the IND-CPA security preservation of the input consistency compiler described in Fig. 14.

valid way, i.e. there exists a random r_i and a message x_i to create a ciphertext ct_i for $i \in [2]$. The compiler is displayed in Fig. 17. In comparison to the NIZK proof system used in the input consistency compiler above, we need to assume one-time simulation-soundness for this advanced case. This leads to the following theorem which is of independent interest beyond the study of consistency.

Theorem 6.6. *Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be an IND-CPA secure functional encryption scheme and $\text{NIZK} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ a NIZK proof system satisfying one-time simulation soundness, then the construction $\text{FE}' = (\text{Setup}', \text{KeyGen}', \text{Enc}', \text{Dec}')$, defined in Figure 17, is IND-CCA secure. Namely, for any PPT adversary \mathcal{A} , there exist PPT adversaries $\mathcal{B}, \mathcal{B}'$ and \mathcal{B}'' , such that:*

$$\text{Adv}_{\text{FE}', \mathcal{A}}^{\text{IND-CCA}}(\lambda) \leq 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{B}}^{\text{ZK}}(\lambda) + \text{Adv}_{\text{NIZK}, \mathcal{B}'}^{\text{Sim-Sound}}(\lambda) + 2 \cdot \text{Adv}_{\text{FE}, \mathcal{B}''}^{\text{IND-CPA}}(\lambda).$$

Proof. To prove this statement, we use a hybrid argument with the games defined in Fig. 19. Note that G_0 corresponds to the game $\text{IND-CCA}_0^{\text{FE}}(1^\lambda, \mathcal{A})$ and G_4 to the game $\text{IND-CCA}_1^{\text{FE}}(1^\lambda, \mathcal{A})$. This results in:

$$\text{Adv}_{\text{FE}', \mathcal{A}}^{\text{IND-CCA}}(1^\lambda) = |\text{Win}_{\mathcal{A}}^{G_0}(1^\lambda) - \text{Win}_{\mathcal{A}}^{G_4}(1^\lambda)| .$$

We describe the different games in more detail:

Game G_1 : In this game, we change from an honestly generated CRS and honestly generated proofs to a simulated CRS and simulated proofs. The transition from G_0 to G_1 is justified by the zero-knowledge property of NIZK. Namely, in Lemma 6.7, we exhibit a PPT adversary \mathcal{B}_0 such that:

$$|\text{Win}_{\mathcal{A}}^{G_0}(1^\lambda) - \text{Win}_{\mathcal{A}}^{G_1}(1^\lambda)| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_0}^{\text{ZK}}(\lambda) .$$

Game G_2 : In this game, we change from an encryption of x^0 to x^1 in the first component of the ciphertext, i.e. $\text{ct} = (\text{Enc}(\text{mpk}_1, x^1), \text{Enc}(\text{mpk}_2, x^0))$. The transition from G_1 to G_2 is justified by the IND-CPA security of FE and the one-time simulation-soundness of NIZK. Namely, in Lemma 6.9, we exhibit PPT adversaries \mathcal{B}_1 and \mathcal{B}_2 such that:

$$|\text{Win}_{\mathcal{A}}^{G_1}(1^\lambda) - \text{Win}_{\mathcal{A}}^{G_2}(1^\lambda)| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_1}^{\text{Sim-Sound}}(\lambda) + \text{Adv}_{\text{FE}, \mathcal{B}_2}^{\text{IND-CPA}}(\lambda) .$$

Game G_3 : In this game, we change from an encryption of x^0 to x^1 in the second component of the ciphertext, i.e. $\text{ct} = (\text{Enc}(\text{mpk}_1, x^1), \text{Enc}(\text{mpk}_2, x^1))$. The transition from G_2 to G_3 is almost symmetric to the transition from game G_1 to G_2 except that it is not necessary to rely on the one-time simulation soundness of the NIZK system and the ciphertext contains an encryption of x^1 in the first position. As in Lemma 6.9,

```

Setup'(1λ) :
CRS ← NIZK.Setup(1λ)
For  $i \in [2]$ 
    ( $\text{mpk}_i, \text{msk}_i$ ) ← Setup(1λ)
Return ( $\text{mpk}', \text{msk}'$ ) = ((CRS,  $\{\text{mpk}_i\}_{i \in [2]}$ ),  $\{\text{msk}_i\}_{i \in [2]}$ )
KeyGen'(mpk', msk', f) :
Parse  $\text{mpk}' = (\text{CRS}, \{\text{mpk}_i\}_{i \in [2]})$ ,  $\text{msk}' = \{\text{msk}_i\}_{i \in [2]}$ 
For  $i \in [2]$ 
     $\text{sk}_{f,i} = \text{KeyGen}(\text{mpk}_i, \text{msk}_i, f)$ 
Return  $\text{sk}'_f = \{\text{sk}_{f,i}\}_{i \in [2]}$ 
Enc'(mpk', x) :
Parse  $\text{mpk}' = (\text{CRS}, \{\text{mpk}_i\}_{i \in [2]})$ 
For  $i \in [2]$ 
     $\text{ct}_i = \text{Enc}(\text{mpk}_i, x; r_i)$  with  $r_i \leftarrow \{0, 1\}^\lambda$ 
If  $\exists i \in [2] : \text{ct}_i = \text{err}$  then return err
Generate  $\pi \leftarrow \text{NIZK.Prove}(\text{CRS}, (\text{mpk}_i, \text{ct}_i)_{i \in [2]}, (x, \{r_i\}_{i \in [2]}))$  for  $R_{\text{in}}^{\text{CCA}}$  (Fig. 18)
If  $\text{NIZK.Verify}(\text{CRS}, (\text{mpk}_i, \text{ct}_i)_{i \in [2]}, \pi) = 0$  return err
Return  $\text{ct}' = (\{\text{ct}_i\}_{i \in [2]}, \pi)$ 
Dec'(mpk', f, sk'_f, ct') :
Parse  $\text{mpk}' := (\text{CRS}, \{\text{mpk}_i\}_{i \in [2]})$ ,  $\text{sk}'_f = \{\text{sk}_{i,f}\}_{i \in [2]}$ ,  $\text{ct}' = (\{\text{ct}_i\}_{i \in [2]}, \pi)$ 
If  $\text{NIZK.Verify}(\text{CRS}, (\text{mpk}_i, \text{ct}_i)_{i \in [2]}, \pi) = 1$ 
    Return  $\text{Dec}(\text{mpk}_1, f, \text{sk}_{f,1}, \text{ct}_1)$ 
Else
    Return  $\perp$ 

```

Fig. 17: Advanced input consistency compiler. Shaded instructions indicate difference to the simpler input consistency compiler.

```

Relation  $R_{\text{in}}^{\text{CCA}}$  :
Instance:  $z = (\text{mpk}_i, \text{ct}_i)_{i \in [2]}$ 
Witness:  $w = (x, \{r_i\}_{i \in [2]})$ ,  $x \in \mathcal{X}$ , random coins  $r_i$ 
 $R_{\text{in}}^{\text{CCA}}(z, w) = 1$  if and only if:
     $\text{ct}_i = \text{Enc}(\text{mpk}_i, x; r_i)$ , for both  $i \in [2]$ 

```

Fig. 18: Relation used in the advanced input consistency compiler.

Game	CRS & π	ct	justification/remark
G_0	$\text{CRS} \leftarrow \text{NIZK.Setup}(1^\lambda)$ $\pi \leftarrow \text{NIZK.Prove}(\text{CRS}, x, w)$	$\text{Enc}(\text{mpk}_1, x^0)$ $\text{Enc}(\text{mpk}_2, x^0)$	
G_1	$\text{CRS} \leftarrow \mathcal{S}_1(1^\lambda)$ $\pi \leftarrow \mathcal{S}_2(\text{CRS}, \tau, x)$	$\text{Enc}(\text{mpk}_1, x^0)$ $\text{Enc}(\text{mpk}_2, x^0)$	Zero-knowledge of NIZK
G_2	$\text{CRS} \leftarrow \mathcal{S}_1(1^\lambda)$ $\pi \leftarrow \mathcal{S}_2(\text{CRS}, \tau, x)$	$\text{Enc}(\text{mpk}_1, x^1)$ $\text{Enc}(\text{mpk}_2, x^0)$	IND-CPA of FE and simulation-soundness of NIZK
G_3	$\text{CRS} \leftarrow \mathcal{S}_1(1^\lambda)$ $\pi \leftarrow \mathcal{S}_2(\text{CRS}, \tau, x)$	$\text{Enc}(\text{mpk}_1, x^1)$ $\text{Enc}(\text{mpk}_2, x^1)$	IND-CPA of FE
G_4	$\text{CRS} \leftarrow \text{NIZK.Setup}(1^\lambda)$ $\pi \leftarrow \text{NIZK.Prove}(\text{CRS}, x, w)$	$\text{Enc}(\text{mpk}_1, x^1)$ $\text{Enc}(\text{mpk}_2, x^1)$	Zero-knowledge of NIZK

Fig. 19: Overview of the games to prove the IND-CCA security of the advanced input consistency compiler described in Fig. 17.

the transition is justified by the IND-CPA security of FE. Namely, we can exhibit a PPT adversary \mathcal{B}_2 such that:

$$|\text{Win}_{\mathcal{A}}^{G_2}(1^\lambda) - \text{Win}_{\mathcal{A}}^{G_3}(1^\lambda)| \leq \text{Adv}_{\text{FE}, \mathcal{B}_1}^{\text{IND-CPA}}(\lambda) .$$

Game G_4 : This game is the $\text{IND-CCA}_1^{\text{FE}}(1^\lambda, \mathcal{A})$ game. The transition from G_3 to G_4 is almost symmetric to the transition from G_0 to G_1 except from the fact that the reduction encrypts x^1 instead of x^0 . As in Lemma 6.7, the transition is justified by the zero-knowledge property of NIZK. Namely, we can exhibit a PPT adversary \mathcal{B}_0 such that:

$$|\text{Win}_{\mathcal{A}}^{G_3}(1^\lambda) - \text{Win}_{\mathcal{A}}^{G_4}(1^\lambda)| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_0}^{\text{ZK}}(\lambda) .$$

Putting everything together, we obtain the theorem. \square

Lemma 6.7 (Transition from G_0 to G_1). *For any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B}_0 such that*

$$|\text{Win}_{\mathcal{A}}^{G_0}(1^\lambda) - \text{Win}_{\mathcal{A}}^{G_1}(1^\lambda)| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_0}^{\text{ZK}}(\lambda) .$$

Proof. We build an adversary \mathcal{B}_0 that simulates G_β towards \mathcal{A} when interacting with the underlying $\text{ZK}_\beta^{\text{NIZK}}$ experiment.

In the beginning of the reduction, \mathcal{B}_0 receives CRS from the $\text{ZK}_\beta^{\text{NIZK}}$ experiment. It generates two functional encryption instances $(\text{mpk}_i, \text{msk}_i) \leftarrow \text{Setup}(1^\lambda)$ for $i \in [2]$, sets $\text{mpk}' = (\text{CRS}, \{\text{mpk}_i\}_{i \in [2]})$ and gives mpk' to the adversary.

Whenever \mathcal{A} asks an encryption query (x^0, x^1) , \mathcal{B}_0 generates the ciphertext $\text{ct}_i = (\text{Enc}(\text{mpk}_i, x^0; r_i))_{i \in [2]}$ with $r_i \leftarrow \{0, 1\}^\lambda$ for $i \in [2]$ and sends $y = (\text{mpk}_i, \text{ct}_i)_{i \in [2]}$ and $w = (x, \{r_i\}_{i \in [2]})$ as a statement-witness pair to its challenger. As an answer, \mathcal{B}_0 receives a proof π for the relation $R_{\text{in}}^{\text{CCA}}$. It sets $\text{ct}' = (\{\text{ct}_i\}_{i \in [2]}, \pi)$ and sends it to \mathcal{A} .

For a key generation query f , \mathcal{B}_0 generates $\text{sk}_{f,i} \leftarrow \text{KeyGen}(\text{mpk}_i, \text{msk}_i, f)$ for $i \in [2]$ and sends $\text{sk}'_f = \{\text{sk}_{f,i}\}_{i \in [2]}$ as a reply to \mathcal{A} .

Whenever \mathcal{A} submits a decryption query (f, ct') , with $\text{ct}' = (\{\text{ct}_i\}_{i \in [2]}, \pi)$, \mathcal{B}_0 generates the functional key $\text{sk}_{f,1} \leftarrow \text{KeyGen}(\text{mpk}_1, \text{msk}_1, f)$ and executes $\text{NIZK.Verify}(\text{CRS}, (\text{mpk}_i, \text{ct}_i)_{i \in [2]})$. If $\text{NIZK.Verify}(\text{CRS}, (\text{mpk}_i, \text{ct}_i)_{i \in [2]}) = 1$, \mathcal{B}_0 computes $y := \text{Dec}(\text{mpk}_1, f, \text{sk}_{f,1}, \text{ct}_1)$ and sends y to \mathcal{A} . Otherwise, \mathcal{B}_0 sends \perp to \mathcal{A} .

This covers the simulation of the game G_β . Finally \mathcal{B}_0 outputs the same bit β' returned by \mathcal{A} . It follows, from the perfect simulation, that the advantage of \mathcal{B}_0 is the same as the advantage of \mathcal{A} . \square

As in [52], we prove a corollary that shows that whenever a decryption oracle query is asked and this query contains a valid NIZK proof, then the corresponding ciphertext is explainable under the queried function. This is necessary for the proof of the transition from G_1 to G_2 for the simulation of the decryption oracle.

Claim 6.8. For any PPT adversary \mathcal{A} participating in $G_{1+\beta}$ for $\beta \in \{0, 1\}$, the probability that, during the experiment, \mathcal{A} queries its decryption oracle QDec with a function-ciphertext-pair that is not explainable but has an accepting proof is negligible. Namely, we exhibit a PPT adversary \mathcal{B}_1 , such that

$$\Pr \left[\begin{array}{l} \exists (f, \{\text{ct}'_i\}_{i \in [2]}, \pi') \in Q : (\{\text{ct}'_i\}_{i \in [2]}, \pi') \neq (\{\text{ct}_i\}_{i \in [2]}, \pi), \\ \text{NIZK.Verify}(\text{CRS}, \{\text{ct}'_i\}_{i \in [2]}, \pi') = 1 \text{ and} \\ \text{Dec}(\text{mpk}_1, \text{sk}_{f,1}, \text{ct}'_1) \neq \text{Dec}(\text{mpk}_2, \text{sk}_{f,2}, \text{ct}'_2) \end{array} \right] \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_1}^{\text{Sim-Sound}}(\lambda),$$

where $\text{sk}_{f,i} \leftarrow \text{KeyGen}(\text{mpk}_i, \text{msk}_i, f)$ for $i \in [2]$, $(\{\text{ct}_i\}_{i \in [2]}, \pi)$ is the reply to the encryption query (x^0, x^1) made by \mathcal{A} and Q the list containing all the decryption queries $(f, \{\text{ct}'_i\}_{i \in [2]}, \pi')$ asked by \mathcal{A} , knowing the master public key $\text{mpk}' := (\text{CRS}, \{\text{mpk}_i\}_{i \in [2]})$, the reply to its challenge query $(\{\text{ct}_i\}_{i \in [2]}, \pi)$ and by having access to the key generation oracle $\text{KeyGen}'(\text{mpk}', \text{msk}', \cdot)$, during the game.

Proof. We build an adversary \mathcal{B}_1 that simulates $G_{1+\beta}$ towards \mathcal{A} when interacting with the underlying one time simulation-soundness experiment.

After the adversary \mathcal{B}_1 has received CRS from the underlying experiment, it generates $(\text{mpk}_i, \text{msk}_i) \leftarrow \text{Setup}(1^\lambda)$ for $i \in [2]$, sets $\text{mpk}' := (\text{CRS}, \{\text{mpk}_i\}_{i \in [2]})$ and sends mpk' to \mathcal{A} . Whenever \mathcal{A} submits a key generation query f , \mathcal{B}_1 generates the functional keys $\text{sk}_{f,i} \leftarrow \text{KeyGen}(\text{mpk}_i, \text{msk}_i, f)$ for $i \in [2]$, sets $\text{sk}'_f = \{\text{sk}_{f,i}\}_{i \in [2]}$ and sends it to \mathcal{A} .

For the challenge query (x^0, x^1) asked by \mathcal{A} , \mathcal{B}_1 computes $\text{ct}_1 = \text{Enc}(\text{mpk}_1, x^\beta)$ and $\text{ct}_2 = \text{Enc}(\text{mpk}_2, x^0)$ (where $\beta = 0$ in game G_1 and $\beta = 1$ in G_2) and asks its experiment for a simulated proof π of the statement $(\text{mpk}_i, \text{ct}_i)_{i \in [2]}$. It sets $\text{ct}' := (\{\text{ct}_i\}_{i \in [2]}, \pi)$ and sends ct' to \mathcal{A} .

Whenever \mathcal{A} outputs a decryption query $(f, \text{ct}' := (\{\text{ct}_i\}_{i \in [2]}, \pi))$, \mathcal{B}_1 verifies the proof. If the output of the verification is 1, \mathcal{B}_1 computes $y_{f,1} = \text{Dec}(\text{mpk}_1, \text{sk}_{f,1}, \text{ct}_1)$ and $y_{f,2} = \text{Dec}(\text{mpk}_2, \text{sk}_{f,2}, \text{ct}_2)$. If $y_{f,1} \neq y_{f,2}$, \mathcal{B}_1 sends $(\{\text{ct}_i\}_{i \in [2]}, \pi)$ as a proof forgery to its challenger. Otherwise it sends $y_{f,1}$ to \mathcal{A} . If the verification outputs 0, \mathcal{B}_1 sends \perp to \mathcal{A} . \square

After introducing and proving Claim 6.8, we prove the transition from G_1 to G_2 .

Lemma 6.9 (Transition from G_1 to G_2). For any PPT adversary \mathcal{A} , there exist PPT adversaries \mathcal{B}_1 and \mathcal{B}_2 , such that

$$|\text{Win}_{\mathcal{A}}^{G_1}(1^\lambda) - \text{Win}_{\mathcal{A}}^{G_2}(1^\lambda)| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_1}^{\text{Sim-Sound}}(\lambda) + \text{Adv}_{\text{FE}, \mathcal{B}_2}^{\text{IND-CPA}}(\lambda) .$$

Proof. We build an adversary \mathcal{B}_2 that simulates $G_{1+\beta}$ to \mathcal{A} when interacting with the underlying IND-CPA $_{\beta}^{\text{FE}}$ experiment.

In the beginning of the reduction, \mathcal{B}_2 receives mpk_1 from the underlying experiment. It simulates a CRS, i.e. $(\text{CRS}, \tau) \leftarrow \mathcal{S}_1(1^\lambda)$, generates a functional encryption instance $(\text{mpk}_2, \text{msk}_2) \leftarrow \text{Setup}(1^\lambda)$ sets $\text{mpk}' := (\text{CRS}, \{\text{mpk}_i\}_{i \in [2]})$ and sends mpk' to \mathcal{A} . Whenever \mathcal{A} submits a key generation query f , \mathcal{B}_2 forwards this query to its own key generation oracle $\text{KeyGen}(\text{mpk}_1, \text{msk}_1, \cdot)$, to receive $\text{sk}_{f,1}$ as an answer. Then, \mathcal{B}_2 generates $\text{sk}_{f,2} \leftarrow \text{KeyGen}(\text{mpk}_2, \text{msk}_2, f)$ by itself, sets $\text{sk}_f = \{\text{sk}_{f,i}\}_{i \in [2]}$ and sends it to \mathcal{A} .

For the challenge query (x^0, x^1) asked by \mathcal{A} , \mathcal{B}_2 forwards it to its own encryption oracle and receives $\text{ct}_1 = \text{Enc}(\text{mpk}_1, x^\beta)$ as an answer. It generates $\text{ct}_2 = \text{Enc}(\text{mpk}_2, x^0)$ by itself, simulates a valid proof π of the relation R_{in} using the statement $y = (\text{mpk}_i, \text{ct}_i)_{i \in [2]}$, i.e. $\pi \leftarrow \mathcal{S}_2(\text{CRS}, \tau, y)$ and sends $\text{ct}' := (\{\text{ct}_i\}_{i \in [2]}, \pi)$ to \mathcal{A} .

Whenever \mathcal{A} asks a decryption query $(f, \text{ct}' := (\{\text{ct}_i\}_{i \in [2]}, \pi))$, \mathcal{B}_2 first verifies the proof π , i.e. it executes $\text{NIZK.Verify}(\text{CRS}, (\text{mpk}_i, \text{ct}_i)_{i \in [2]}, \pi)$. If the verification outputs 1, \mathcal{B}_2 generates $\text{sk}_{f,2} \leftarrow \text{KeyGen}(\text{mpk}_2, \text{msk}_2, f)$, executes $\text{Dec}(\text{mpk}_2, f, \text{sk}_{f,2}, \text{ct}_2)$ and sends the result to \mathcal{A} . This is contrary to the actual decryption oracle, which would always generate the key $\text{sk}_{f,1} \leftarrow \text{KeyGen}(\text{mpk}_1, \text{msk}_1, f)$ and use it to decrypt the first ciphertext ct_1 . Since Claim 6.8 shows that for all ciphertext queries made by \mathcal{A} that have a valid proof (except with negligible probability) it holds that $\text{Dec}(\text{mpk}_1, \text{sk}_{f,1}, \text{ct}_1) = \text{Dec}(\text{mpk}_2, \text{sk}_{f,2}, \text{ct}_2)$. Therefore, it is possible to generate a functional key for either of the position and use the corresponding decryption output as a reply for \mathcal{A} ¹². If the verification outputs 0, \mathcal{B}_2 sends \perp to \mathcal{A} .

This covers the simulation of the game $\text{G}_{1+\beta}$. Finally \mathcal{B}_2 outputs the same bit β' returned by \mathcal{A} . Together with the analysis of adversary \mathcal{B}_1 , this yields the advantage mentioned in the lemma. \square

After proving that the compiler does the security lifting from CPA to CCA, we also need to show that the compiler achieves input consistency.

Theorem 6.10. *Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be a functional encryption scheme and $\text{NIZK} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ a NIZK proof system for $R_{\text{set}}^{\text{CCA}}$ (Fig. 24), then the construction $\text{FE}' = (\text{Setup}', \text{KeyGen}', \text{Enc}', \text{Dec}')$ defined in Figure 17 is input consistent. Namely, for any PPT adversary \mathcal{A} , exists a PPT adversary \mathcal{B} such that*

$$|\Pr[\text{in-CONS}^{\text{FE}'}(1^\lambda, \mathcal{A}, \text{Extract}) = 1]| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}}^{\text{Sound}}(\lambda),$$

where Extract is the algorithm that always returns unknown.

Proof. The proof proceeds exactly in the same way as the proof of Theorem 6.1. It is ensured by the soundness of the NIZK proof that both of the ciphertexts encrypt the same underlying message. \square

Instantiations Our compilers can be instantiated with any NIZK scheme such as [46, 47]. An important special case are lattice based constructions. Since recent results [57, 60] show how to construct NIZKs from LWE, it can be combined with a functional encryption scheme for all classes of circuits from LWE known from [28, 43], to obtain a specific instantiation of the presented compiler from LWE. Furthermore, specific versions of the compiler for restricted classes, namely inner-product classes achievable from standard assumptions [8], could be obtained by employing designated verifier NIZKs from [27, 31, 49, 59]. Since the latter are based on the Diffie-Hellman assumption, the overall scheme is input-consistent based on standard assumptions. We note that it is an interesting research direction to investigate efficient constructions of consistent FE schemes.

6.2 Setup Consistency Compilers

First Compiler. To achieve setup consistency, we need to prevent the generation of malicious functional keys under maliciously generated parameters. While we can still rely on honest encryption procedures, the parameters are chosen by the adversary beforehand and we cannot rely on a common-reference string generated by the adversary.

We replace the role of the NIZK proof in the previous section by a non-interactive witness indistinguishable (NIWI) proof. NIWI proofs allow us to achieve similar properties in terms of correctness and soundness, as provided by the NIZK proof, without relying on a common reference string. As a trade-off, we cannot rely on the zero-knowledge property but on witness-indistinguishability instead, which we prove to be sufficient. However, our compiler needs to run three different instances of the same functional encryption scheme in parallel. The decryption procedure then computes the decryption under all of the three instances and outputs the majority of the decryptions. If no majority is reached, the algorithm outputs \diamond . We give a formal description of this compiler in Fig. 20. This compiler only preserves CPA and, under certain conditions, also CFE security. Afterwards we also present a compiler that achieves CCA security by relying on a CPA secure scheme.

We start by proving the setup consistency of the compiler.

¹² It is not necessary to rely on Claim 6.8 in the transition from G_2 to G_3 , since \mathcal{B}_2 is able to generate a functional key for the first position and therefore is able to generate the decryption oracle perfectly.

Theorem 6.11. *Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be a functional encryption scheme and $\text{NIWI} = (\text{NIWI.Prove}, \text{NIWI.Verify})$ a NIWI proof system for R_{set} (Fig. 21), then the construction $\text{FE}' = (\text{Setup}', \text{KeyGen}', \text{Enc}', \text{Dec}')$ defined in Figure 20 is setup consistent. Namely, for any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B} such that:*

$$|\Pr[\text{set-CONS}^{\text{FE}'}(1^\lambda, \mathcal{A}) = 1]| \leq \text{Adv}_{\text{NIWI}, \mathcal{B}}^{\text{Sound}}(\lambda).$$

Proof. To prove the setup consistency of the scheme FE' , we rely on the soundness of the NIWI proof system. In more detail, we construct an adversary \mathcal{B} that successfully generates a valid proof for a statement not in the language by assuming an adversary \mathcal{A} for the setup consistency experiment set-CONS .

In the beginning of the reduction, \mathcal{B} receives $(\text{mpk}' := (\text{mpk}_1, \text{mpk}_2, \text{mpk}_3), \text{mpk}'' := (\text{mpk}'_1, \text{mpk}'_2, \text{mpk}'_3), \text{sk} := (\{\text{sk}_i\}_{i \in [3]}, \pi), f, x_1, x_2)$ from \mathcal{A} . If π is such that the condition $\text{NIWI.Verify}(1^\lambda, (\{\text{mpk}''_i\}_{i \in [3]}, \{\text{sk}_i\}_{i \in [3]}, f), \pi) = 0$ or $\text{mpk}'' \neq \text{mpk}'$, then \mathcal{B} halts. Note that in this case, \mathcal{A} would never win, as the outcome of decryption procedure is equal to “ \diamond ” for all ciphertexts. Another case, in which \mathcal{B} halts, is the case in which $\text{Dec}'(\text{mpk}'', f, \text{sk}, \text{Enc}(\text{mpk}', x_i)) = f(x_i)$ for both $i \in [2]$. Here, the generated functional key and the public parameters have an honest behavior and therefore \mathcal{A} has not generated a forgery for the NIWI proof. Therefore the adversary \mathcal{B} halts. Otherwise, the adversary outputs the statement $(\{\text{mpk}''_i\}_{i \in [3]}, \{\text{sk}_i\}_{i \in [3]}, f)$ and proof π as a NIWI forgery.

Let us analyze the output of \mathcal{A} to see that the condition to break setup consistency must imply a soundness violation of the NIWI scheme. In more detail, we define the event E as the event that the adversary \mathcal{A} performs a consistency attack under the assumption that $(\{\text{mpk}''_i\}_{i \in [3]}, \{\text{sk}_i\}_{i \in [3]}, f) \in L$ and show that the occurrence of E would contradict the assumption. Let us compute $\text{ct}'_i := (\text{mpk}', \{\text{ct}_j\}_{j \in [3]}) \leftarrow \text{Enc}'(\text{mpk}', x_i)$ for $i \in [2]$ and $y_i = \text{Dec}'(\text{mpk}'', f, \text{sk}, \text{ct}'_i)$ for $i \in [2]$. For concreteness, assume that both ciphertexts are not equal to err (however, the argument holds for any pattern, since erroneous ciphertexts are ignored in the setup consistency game).

Now, we analyze the possible outcomes for the decryptions y_1 and y_2 in the case of a consistency attack. We show that $y_i \neq \diamond$, we denote this by event E_1 , enforces that $y_i = f(x_i)$ for all $i \in [2]$ and furthermore that $y_i \neq \diamond$ for all $i \in [2]$.

In the case of event E_1 , we assume $y_1 \neq \diamond$ (we do the analysis for y_1 , the case for y_2 follows respectively) and, for the sake of contradiction, we also need to assume that $(\{\text{mpk}''_i\}_{i \in [3]}, \{\text{sk}_i\}_{i \in [3]}, f) \in L$ holds. Under these circumstances consistency must be satisfied.

By the perfect correctness of the underlying FE scheme and the validity of the proof, at least two functional keys sk_i and sk_j , for $i \neq j$ are correctly generated and matching to the master public keys mpk_i and mpk_j in the encryption (note that by the definition of Dec' that decryption is only performed if both Enc' and Dec' use the same triple $(\text{mpk}_1, \text{mpk}_2, \text{mpk}_3)$) and thus $y_{f,1}^{(i)} \leftarrow \text{Dec}(\text{mpk}_i, f, \text{sk}_i, \text{ct}_1)$ and $y_{f,1}^{(j)} \leftarrow \text{Dec}(\text{mpk}_j, f, \text{sk}_j, \text{ct}_1)$ are equal to $f(x_1)$. Therefore also the majority of the decryption values for ct_1 is equal to $f(x_1)$ and the final decryption outputs $f(x_1)$. Hence assuming $y_1 \neq \diamond$ implies $y_1 = f(x_1)$.

For event E_2 , we need to show that $y_1 \neq \diamond$ and $y_2 \neq \diamond$ in the case of a consistency attack and under the assumption that $(\{\text{mpk}''_i\}_{i \in [3]}, \{\text{sk}_i\}_{i \in [3]}, f) \in L$. We start by considering the case that $y_1 = y_2 = \diamond$. If $y_1 = y_2 = \diamond$, then the adversary \mathcal{A} did not perform a consistency attack. This is a contradiction to our assumption and therefore this case cannot occur. In the next step we assume that $y_1 \neq \diamond$ and $y_2 = \diamond$ (or $y_1 = \diamond$ and $y_2 \neq \diamond$ respectively). If $y_1 \neq \diamond$, then follows, with the analysis for E_1 , that $y_1 = f(x_1)$ and that at least two of the functional keys sk_i and sk_j are correctly generated and matching the master public keys mpk_i and mpk_j . But this would also lead, due to perfect correctness of the functional encryption scheme, to a correct decryption of the ciphertext ct'_2 , which yields $y_2 = f(x_2)$. This shows that the case $y_1 \neq \diamond$ and $y_2 = \diamond$ (or $y_1 = \diamond$ and $y_2 \neq \diamond$) cannot occur.

By combining the events E_1 and E_2 , we proved that event E cannot occur. To recap, whenever $(\{\text{mpk}''_i\}_{i \in [3]}, \{\text{sk}_i\}_{i \in [3]}, f) \in L$, it is not possible for an adversary \mathcal{A} to perform a consistency attack. Hence the only way the adversary can break setup consistency is by breaking the soundness property of the NIWI scheme, i.e., providing the statement $(\{\text{mpk}''_i\}_{i \in [3]}, \{\text{sk}_i\}_{i \in [3]}, f) \notin L$.

This yields the bound

$$|\Pr[\text{set-CONS}^{\text{FE}'}(1^\lambda, \mathcal{A}) = 1]| \leq \text{Adv}_{\text{NIWI}, \mathcal{B}}^{\text{Sound}}(\lambda) .$$

and therefore we obtain the theorem. \square

Besides proving that the compiler achieves setup consistency, we need to show that the security of the underlying functional encryption scheme is preserved.

<p><u>Setup'(1$^\lambda$) :</u> For $i \in [3]$: $(\text{mpk}_i, \text{msk}_i) \leftarrow \text{Setup}(1^\lambda; s_i)$ $\text{mpk}' := \{\text{mpk}_i\}_{i \in [3]}$ $\text{msk}' := \{(\text{msk}_i, s_i)\}_{i \in [3]}$ Return $(\text{mpk}', \text{msk}')$</p> <p><u>KeyGen'(mpk', msk', f) :</u> Parse $\text{mpk}' := \{\text{mpk}_i\}_{i \in [3]}$, $\text{msk}' := \{(\text{msk}_i, s_i)\}_{i \in [3]}$ For $i \in [3]$: $r_i \in \{0, 1\}^\lambda$ $\text{sk}_{f,i} = \text{KeyGen}(\text{msk}_i, f; r_i)$ Generate $\pi \leftarrow \text{NIWI.Prove}(1^\lambda, z, w)$ with $z = (\{\text{mpk}_i\}_{i \in [3]}, \{\text{sk}_{f,i}\}_{i \in [3]}, f)$, $w = (\{\text{msk}_i\}_{i \in [3]}, \{r_i\}_{i \in [3]}, \{s_i\}_{i \in [3]})$ and L defined corresponding to R_{set} (Fig. 21) Return $\text{sk}'_f = (\{\text{sk}_{f,i}\}_{i \in [3]}, \pi)$</p>	<p><u>Enc'(mpk', x) :</u> Parse $\text{mpk}' := \{\text{mpk}_i\}_{i \in [3]}$ For $i \in [3]$: $\text{ct}_i \leftarrow \text{Enc}(\text{mpk}_i, x)$ If $\exists i \in [3] : \text{ct}_i = \text{err}$ then return err Return $\text{ct}' = (\text{mpk}', \{\text{ct}_i\}_{i \in [3]})$</p> <p><u>Dec'(mpk', f, sk'_f, ct') :</u> Parse $\text{mpk}' := \{\text{mpk}_i\}_{i \in [3]}$, $\text{sk}'_f = (\{\text{sk}_{f,i}\}_{i \in [3]}, \pi)$, $\text{ct}' = (\text{mpk}'', \{\text{ct}_i\}_{i \in [3]})$ $z := (\{\text{mpk}_i\}_{i \in [3]}, \{\text{sk}_{f,i}\}_{i \in [3]}, f)$ If $\text{mpk}' = \text{mpk}''$ If $\text{NIWI.Verify}(1^\lambda, z, \pi) = 1$ For $i \in [3]$: $y_{f,i} := \text{Dec}(\text{mpk}_i, f, \text{sk}'_{f,i}, \text{ct}_i)$, $y \leftarrow \text{MajVal}(\{y_{f,i}\}_{i \in [3]})$ Return y Return \diamond</p>
---	---

Fig. 20: Setup consistency compiler. $\text{MajVal}(\cdot)$ calculates and returns the majority value of the input values, if there is a clear majority and \diamond otherwise.

CPA case. We start with the CPA case, which is straightforward:

Theorem 6.12. *Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be an IND-CPA secure functional encryption scheme and $\text{NIWI} = (\text{NIWI.Prove}, \text{NIWI.Verify})$ a NIWI proof system for R_{set} (Fig. 21), then the construction FE' defined in Figure 20 is IND-CPA secure. Namely, for any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B} such that:*

$$\text{Adv}_{\text{FE}', \mathcal{A}}^{\text{IND-CPA}}(\lambda) \leq 3 \cdot \text{Adv}_{\text{FE}, \mathcal{B}}^{\text{IND-CPA}}(\lambda) .$$

Proof. To prove this statement, we use a hybrid argument with the games $\text{G}_0, \dots, \text{G}_3$ as defined in Fig. 22. Note that G_0 corresponds to the game $\text{IND-CPA}_0^{\text{FE}}(1^\lambda, \mathcal{A})$ and G_3 to the game $\text{IND-CPA}_1^{\text{FE}}(1^\lambda, \mathcal{A})$. By using the triangle inequality, we get:

$$\text{Adv}_{\text{FE}', \mathcal{A}}^{\text{IND-CPA}}(\lambda) \leq \sum_{i=k}^3 |\text{Win}_{\mathcal{A}}^{\text{G}_{k-1}}(1^\lambda) - \text{Win}_{\mathcal{A}}^{\text{G}_k}(1^\lambda)| .$$

We conclude the proof by showing that for any $k \in [3]$, there exists an adversary \mathcal{B}_k , such that:

<p>Relation R_{set} :</p> <p>Instance: $z = (\{\text{mpk}_i\}_{i \in [3]}, \{\text{sk}_{f,i}\}_{i \in [3]}, f)$</p> <p>Witness: $w = (\{\text{msk}_i\}_{i \in [3]}, \{r_i\}_{i \in [3]}, \{s_i\}_{i \in [3]})$</p> <p>$R_{\text{set}}(z, w) = 1$ if and only if for at least two indices $j_1, j_2 \in [3]$, $j_1 \neq j_2$ we have:</p> <ol style="list-style-type: none"> The key sk_{f,j_1} for function f is generated using msk_{j_1} with r_{j_1} and related to mpk_{j_1} which is generated using s_{j_1}. Formally: $\text{sk}_{f,j_1} = \text{KeyGen}(\text{mpk}_{j_1}, \text{msk}_{j_1}, f; r_{j_1}) \wedge (\text{mpk}_{j_1}, \text{msk}_{j_1}) = \text{Setup}(1^\lambda; s_{j_1})$ (AND) The key sk_{f,j_2} for function f is generated using msk_{j_2} with r_{j_2} and related to mpk_{j_2} which is generated using s_{j_2}. Formally: $\text{sk}_{f,j_2} = \text{KeyGen}(\text{mpk}_{j_2}, \text{msk}_{j_2}, f; r_{j_2}) \wedge (\text{mpk}_{j_2}, \text{msk}_{j_2}) = \text{Setup}(1^\lambda; s_{j_2})$

Fig. 21: Relation used in the setup consistency compiler

Game	ct	justification/remark
G_k	For $i \leq k$ $\text{Enc}(\text{mpk}_i, x^0)$ For $i > k$ $\text{Enc}(\text{mpk}_i, \boxed{x^1})$	IND-CPA security of FE

Fig. 22: Overview of the games to prove the IND-CPA security preservation of the setup consistency compiler described in Fig. 20.

$$|\text{Win}_{\mathcal{A}}^{G_{k-1}}(1^\lambda) - \text{Win}_{\mathcal{A}}^{G_k}(1^\lambda)| \leq \text{Adv}_{\text{FE}, \mathcal{B}_k}^{\text{IND-CPA}}(\lambda) .$$

The adversary \mathcal{B} of the statement is defined as the monolithic adversary that first samples $k \leftarrow [s]$ uniformly at random and then runs the code of \mathcal{B}_k .

We build an adversary \mathcal{B}_k that simulates $G_{k+\beta}$ towards \mathcal{A} when interacting with the underlying $\text{IND-CPA}_{\beta}^{\text{FE}}$ experiment.

In the beginning of the reduction, \mathcal{B}_k receives mpk_k from the experiment. It generates two other functional encryption instances $(\text{mpk}_i, \text{msk}_i) \leftarrow \text{Setup}(1^\lambda)$, for $i \in [3] \setminus \{k\}$, sets $\text{mpk}' = \{\text{mpk}_i\}_{i \in [3]}$ and gives mpk' to the adversary.

Whenever \mathcal{A} asks an encryption query (x^0, x^1) , \mathcal{B}_k forwards it to its own encryption oracle to receive $\text{ct}_k \leftarrow \text{Enc}(\text{mpk}_k, x^\beta)$, generates $\text{ct}_i \leftarrow \text{Enc}(\text{mpk}_i, x^1)$, for $i < k$, and $\text{ct}_i \leftarrow \text{Enc}(\text{mpk}_i, x^0)$, for $i > k$, on its own and sends $\text{ct}' = \{\text{ct}_i\}_{i \in [3]}$ to \mathcal{A} .

For a key generation query f , \mathcal{B}_k queries its own key generation oracle on f to receive $\text{sk}_{f,k} \leftarrow \text{KeyGen}(\text{mpk}_k, \text{msk}_k, f)$, generates $\text{sk}_{f,i} \leftarrow \text{KeyGen}(\text{mpk}_i, \text{msk}_i, f)$ for $i \in [3] \setminus \{k\}$ on its own and generates a proof $\pi \leftarrow \text{NIWI.Prove}(1^\lambda, z, w)$ with $z = (\{\text{mpk}_i\}_{i \in [3]}, \{\text{sk}_{f,i}\}_{i \in [3]}, f)$, $w = (\{\text{msk}_i\}_{i \in [3] \setminus \{k\}}, \{r_i\}_{i \in [3] \setminus \{k\}}, \{s_i\}_{i \in [3] \setminus \{k\}})$ for the relation R_{set} , by using its information of two-out-of-the-three different instances. As a reply for the key generation query, \mathcal{B}_k sends $\text{sk}'_f := (\{\text{sk}_{f,i}\}_{i \in [3]}, \pi)$ to \mathcal{A} .

This covers the simulation of the game $G_{k+\beta}$. Finally \mathcal{B}_k outputs the same bit β' returned by \mathcal{A} . It follows, from the perfect simulation, that the advantage of \mathcal{B}_k is the same as the advantage \mathcal{A} . \square

CFE case. Due to the much stronger simulation-based security requirement of CFE, the existence of simulators \mathcal{S}_1 (for setup generation), that outputs a simulated mpk and an initial (joint) state $s = s_{\text{init}}$, \mathcal{S}_2 for the simulation of functional keys (based on the joint state s which might be update in this process), and finally

\mathcal{S}_3 for simulating ciphertexts (again with access to the joint state s) does formally not imply knowledge of a master secret key that would be needed to create valid proofs for the relation R_{set} . Hence, the theorem captures preservation only for a specific class of simulators and not all CFE secure schemes. We note that for the brute-force scheme in Fig. 12 in Section 4 there exists a simulator that belongs to the class we are proving the security preservation for.

More formally, we require \mathcal{S}_1 to output mpk and maintain state s such that $(\text{mpk}, \text{msk}) = \text{Setup}(1^\lambda; s)$ holds with probability 1 in $\text{Ideal}^{\text{FE}}(1^\lambda, \mathcal{A}, \mathcal{S})$. Additionally, we require that for any adversary \mathcal{A} any functional key sk_f output by \mathcal{S}_2 on input f satisfies $\text{sk}_f = \text{KeyGen}(\text{mpk}, \text{msk}, f)$ with probability 1 in $\text{Ideal}^{\text{FE}}(1^\lambda, \mathcal{A}, \mathcal{S})$ where mpk and msk are the values obtained by $\text{Setup}(1^\lambda; s)$ where s is the initial private state output by \mathcal{S}_1 .

Note that the simulator for the brute-force scheme described in [53] runs the normal setup-algorithm in the simulation (and can hence provide the randomness used during the generation) and the master secret key fixes all secret keys.

Theorem 6.13. *Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be a CFE secure functional encryption scheme with respect to simulators $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3)$ that satisfy the above condition in $\text{Ideal}^{\text{FE}}(1^\lambda, \mathcal{A}, \mathcal{S})$ (for any adversary \mathcal{A}). Let further $\text{NIWI} = (\text{NIWI.Prove}, \text{NIWI.Verify})$ be a NIWI proof system for R_{set} (Fig. 21). Under the assumption that KeyGen is deterministic, the construction FE' defined in Figure 20 is CFE secure.*

Proof. Under the theorem's assumptions, the simulator \mathcal{S}_1 's output is essentially equivalent to the master secret key. Together with the fact that key derivation is deterministic, we see that all NIWI proofs can be simulated. More detailed, we can run three independent simulations of the FE scheme for the overall simulation. That is, let $\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3$ be the simulators for FE. Then the composite simulator $\mathcal{S}' = (\mathcal{S}'_1, \mathcal{S}'_2, \mathcal{S}'_3)$ works as follows: \mathcal{S}'_1 runs \mathcal{S}_1 three times to obtain (mpk_i, s_i) . To answer key-generation queries for functions f , \mathcal{S}'_2 runs \mathcal{S}_2 three times on the respective joint state (and the function values of all previous queries) to obtain $\text{sk}_{f,i}$. Note that by the assumption on the simulation for the underlying scheme, for $(\text{mpk}_i, \text{msk}_i) = \text{Setup}(1^\lambda; s_i)$ we have that $\text{sk}_{f,i} = \text{KeyGen}(\text{mpk}_i, \text{msk}_i, f)$. By the theorem assumption, KeyGen is deterministic and hence we have all witnesses to simulate a genuine proof $\pi \leftarrow \text{NIWI.Prove}(1^\lambda, z, w)$ with $z = (\{\text{mpk}_i\}_{i \in [3]}, \{\text{sk}_{f,i}\}_{i \in [3]}, f)$ and $w = (\{s_i\}_{i \in [3]}, \{r_i\}_{i \in [3]})$. Finally, simulating a ciphertext is done by invoking \mathcal{S}_3 three times on all three simulated instances (and on the joint state and the function values of the actual plaintext). \square

Second Advanced Compiler. For the advanced setup consistency compiler that takes a CPA secure scheme and achieves CCA security, we proceed in a similar way as in the input consistency case. This is possible since security (in the sense of confidentiality) is only required w.r.t. an honest setup generator. Therefore, as long as the stronger tools required by the Naor-Yung approach [55] do smoothly integrate and not interfere with the tools needed to obtain setup consistency as of Theorem 6.11, we can follow a similar path, but have to pay attention to the details regarding the interplay of the three FE instances. Due to the fact that we already run three instances of the functional encryption scheme, we do not need to run any additional instance, but we need to generate a NIZK proof for every generated encryption. In more detail, we create a NIZK proof that shows that all of the three encryptions are generated in a valid way, i.e. there exists a random r_i and a message x_i to create the ciphertext ct_i for all $i \in [3]$. The compiler is displayed in Fig. 23.

The NIZK system here also needs to fulfill the notion of one-time simulation-soundness in this advanced case. This leads to the following theorem:

Theorem 6.14. *Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be an IND-CPA secure functional encryption scheme, $\text{NIWI} = (\text{NIWI.Prove}, \text{NIWI.Verify})$ a NIWI proof system and $\text{NIZK} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ a NIZK proof system system satisfying one-time simulation soundness, then the construction FE' defined in Figure 23 is IND-CCA secure. Namely, for any PPT adversary \mathcal{A} , there exist PPT adversaries $\mathcal{B}, \mathcal{B}'$ and \mathcal{B}'' such that:*

$$\text{Adv}_{\text{FE}', \mathcal{A}}^{\text{IND-CCA}}(\lambda) \leq 2 \cdot \text{Adv}_{\text{NIZK}, \mathcal{B}'}^{\text{ZK}}(\lambda) + 3 \cdot \text{Adv}_{\text{NIZK}, \mathcal{B}'}^{\text{Sim-Sound}}(\lambda) + 3 \cdot \text{Adv}_{\text{FE}, \mathcal{B}''}^{\text{IND-CPA}}(\lambda).$$

Proof. To prove this statement, we use a hybrid argument with the games defined in Fig. 25. Note that \mathcal{G}_0 corresponds to the game $\text{IND-CCA}_0^{\text{FE}}(1^\lambda, \mathcal{A})$ and \mathcal{G}_3 to the game $\text{IND-CCA}_1^{\text{FE}}(1^\lambda, \mathcal{A})$. This results in:

```

Setup'(1λ) :
  CRS ← NIZK.Setup(1λ)
  For i ∈ [3]:
    (mpki, mski) ← Setup(1λ; si) with si ← {0, 1}λ
  Return (mpk', msk') = (({mpki}i∈[3], CRS), {(mski, si)i∈[3])
KeyGen'(mpk', msk', f) :
  Parse mpk' := ({mpki}i∈[3], CRS), msk' := {(mski, si)i∈[3]
  For i ∈ [3]:
    skf,i = KeyGen(mpki, mski, f; ri) with ri ∈ {0, 1}λ
  Generate πsk ← NIWI.Prove(1λ, z, w) with
  z = ({mpki}i∈[3], {skf,i}i∈[3], f)
  w = ({mski}i∈[3], {ri}i∈[3], {si}i∈[3])
  where L is defined corresponding to Rset (Fig. 21)
  Return sk'_f = ({skf,i}i∈[3], πsk)
Enc'(mpk', x) :
  Parse mpk' := ({mpki}i∈[3], CRS)
  For i ∈ [3]:
    cti ← Enc(mpki, x; ui) with ui ∈ {0, 1}λ
  If ∃ i ∈ [3] : cti = err return err
  Generate πct ← NIZK.Prove(CRS, (mpki, cti)i∈[3], (x, {ui}i∈[3])),
  for RsetCCA (Fig. 24)
  If NIZK.Verify(CRS, (mpki, cti)i∈[3], πct) = 0 return err
  Return ct' = (mpk', {cti}i∈[3], πct)
Dec'(mpk', sk'_f, f, ct') :
  Parse mpk' := {mpki}i∈[3], sk'_f = ({skf,i}i∈[3], πsk), ct' = (mpk'', {cti}i∈[3], πct)
  If mpk' = mpk''
    If NIZK.Verify(CRS, (mpki, cti)i∈[3], πct) = 1
      If NIWI.Verify(1λ, ({mpki}i∈[3], {skf,i}i∈[3], f), πsk) = 1
        yf,i := Dec(mpki, skf,i, f, cti), for i ∈ [3]:
        If there are indices a, b ∈ [3], a ≠ b s.t. yf,a = yf,b
          Return y ← MajVal(yf,1, yf,2, yf,3)
  Return ◊

```

Fig. 23: Advanced setup consistency compiler. MajVal(·) calculates and returns the majority value of the input values, if there is a clear majority and ◊ otherwise. Shaded instructions again indicate the difference to the simpler setup compiler.

<p>Relation $R_{\text{set}}^{\text{CCA}}$:</p> <p>Instance: $z = (\{\text{mpk}_i\}_{i \in [3]}, \{\text{ct}_i\}_{i \in [3]})$</p> <p>Witness: $w = (x, \{u_i\}_{i \in [3]})$</p> <p>$R_{\text{set}}^{\text{CCA}}(z, w) = 1$ if and only if:</p> <p style="text-align: center;">$\text{ct}_i = \text{Enc}(\text{mpk}_i, x; u_i)$ for all $i \in [3]$</p>
--

Fig. 24: Relation used in the advanced setup consistency compiler

Game	CRS & π	ct	justification/remark
G_0	$\text{CRS} \leftarrow \text{NIZK.Setup}(1^\lambda)$ $\pi \leftarrow \text{NIZK.Prove}(\text{CRS}, x, w)$	For $i \in [3]$ $\text{Enc}(\text{mpk}_i, x^0)$	
G_1	$\text{CRS} \leftarrow \mathcal{S}_1(1^\lambda)$ $\pi \leftarrow \mathcal{S}_2(\text{CRS}, \tau, x)$	For $i \in [3]$ $\text{Enc}(\text{mpk}_i, x^0)$	Zero-knowledge of NIZK
$G_{1.k}$	$\text{CRS} \leftarrow \mathcal{S}_1(1^\lambda)$ $\pi \leftarrow \mathcal{S}_2(\text{CRS}, \tau, x)$	For $i \leq k$ $\text{Enc}(\text{mpk}_i, \boxed{x^1})$ For $i > k$ $\text{Enc}(\text{mpk}_i, x^0)$	IND-CPA of FE and one-time simulation-soundness of NIZK
G_2	$\text{CRS} \leftarrow \mathcal{S}_1(1^\lambda)$ $\pi \leftarrow \mathcal{S}_2(\text{CRS}, \tau, x)$	For $i \in [3]$ $\text{Enc}(\text{mpk}_i, x^1)$	$G_2 = G_{1.3}$
G_3	$\text{CRS} \leftarrow \text{NIZK.Setup}(1^\lambda)$ $\pi \leftarrow \text{NIZK.Prove}(\text{CRS}, x, w)$	For $i \in [3]$ $\text{Enc}(\text{mpk}_i, x^0)$	Zero-knowledge of NIZK

Fig. 25: Overview of the games to prove the IND-CCA security of the advanced setup consistency compiler described in Fig. 23.

$$\text{Adv}_{\text{FE}', \mathcal{A}}^{\text{IND-CCA}}(1^\lambda) = |\text{Win}_{\mathcal{A}}^{\text{G}_0}(1^\lambda) - \text{Win}_{\mathcal{A}}^{\text{G}_3}(1^\lambda)| .$$

We describe the different games in more detail:

Game G_1 : In this game, we change from an honestly generated CRS and honestly generated proofs to a simulated CRS and simulated proofs. The transition from G_0 to G_1 is justified by the zero-knowledge property of NIZK. Namely, in Lemma 6.15, we exhibit a PPT adversary \mathcal{B}_0 such that:

$$|\text{Win}_{\mathcal{A}}^{\text{G}_0}(1^\lambda) - \text{Win}_{\mathcal{A}}^{\text{G}_1}(1^\lambda)| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_0}^{\text{ZK}}(\lambda) .$$

Game G_2 : In this game, we change from an encryption of x^0 to x^1 in the three components of the ciphertext. The transition from G_1 to G_2 is justified by the IND-CPA security of FE and the one-time simulation-soundness of NIZK. Namely, in Lemma 6.17, we exhibit PPT adversaries \mathcal{B}_1 and \mathcal{B}_2 such that:

$$|\text{Win}_{\mathcal{A}}^{\text{G}_1}(1^\lambda) - \text{Win}_{\mathcal{A}}^{\text{G}_2}(1^\lambda)| \leq 3 \cdot \text{Adv}_{\text{NIZK}, \mathcal{B}_1}^{\text{Sim-Sound}}(\lambda) + 3 \cdot \text{Adv}_{\text{FE}, \mathcal{B}_2}^{\text{IND-CPA}}(\lambda) .$$

Game G_3 : This game is the $\text{IND-CCA}_1^{\text{FE}}(1^\lambda, \mathcal{A})$ game. The transition from G_4 to G_5 is almost symmetric to the transition from G_0 to G_1 except from the fact that the reduction encrypts x^1 instead of x^0 . As in Lemma 6.15, the transition is justified by the zero-knowledge property of NIZK. Namely, we can exhibit a PPT adversary \mathcal{B}_0 such that:

$$|\text{Win}_{\mathcal{A}}^{\text{G}_3}(1^\lambda) - \text{Win}_{\mathcal{A}}^{\text{G}_4}(1^\lambda)| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_0}^{\text{ZK}}(\lambda) .$$

Putting everything together, we obtain the theorem. □

Lemma 6.15 (Transition from G_0 to G_1). *For any PPT adversary \mathcal{A} , there exists a PPT adversary \mathcal{B}_0 such that*

$$|\text{Win}_{\mathcal{A}}^{\text{G}_0}(1^\lambda) - \text{Win}_{\mathcal{A}}^{\text{G}_1}(1^\lambda)| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_0}^{\text{ZK}}(\lambda) .$$

Proof. We build an adversary \mathcal{B}_0 that simulates G_β towards \mathcal{A} when interacting with the underlying $\text{ZK}_\beta^{\text{NIZK}}$ experiment.

In the beginning of the reduction, \mathcal{B}_0 receives CRS from the $\text{ZK}_\beta^{\text{NIZK}}$ experiment. It generates three functional encryption instances $(\text{mpk}_i, \text{msk}_i) \leftarrow \text{Setup}(1^\lambda; s_i)$ with $s_i \leftarrow \{0, 1\}^\lambda$ for $i \in [3]$, sets $\text{mpk}' = (\text{CRS}, \{\text{mpk}_i\}_{i \in [3]})$ and gives mpk' to the adversary.

Whenever \mathcal{A} asks an encryption query (x^0, x^1) , \mathcal{B}_0 generates the ciphertext $\text{ct}_i = (\text{Enc}(\text{mpk}_i, x^0; u_i))_{i \in [3]}$ with $u_i \leftarrow \{0, 1\}^\lambda$ for $i \in [3]$ and sends $y = (\text{mpk}_i, \text{ct}_i)_{i \in [3]}$ and $w = (x, \{u_i\}_{i \in [3]})$ as a statement-witness pair to its challenger. As an answer, \mathcal{B}_0 receives a proof π_{ct} for the relation $R_{\text{set}}^{\text{CCA}}$. It sets $\text{ct}' = (\{\text{ct}_i\}_{i \in [3]}, \pi_{\text{ct}})$ and sends it to \mathcal{A} .

For a key generation query f , \mathcal{B}_0 generates $\text{sk}_{f,i} \leftarrow \text{KeyGen}(\text{mpk}_i, \text{msk}_i, f; r_i)$ with $r_i \leftarrow \{0, 1\}^\lambda$ for $i \in [3]$ and creates a NIWI proof π_{sk} over the relation R_{set} for the statement-witness pair $y = (\{\text{mpk}_i\}_{i \in [3]}, \{\text{sk}_{f,i}\}_{i \in [3]}, f)$ and $w = (\{\text{msk}_i\}_{i \in [3]}, \{r_i\}_{i \in [3]}, \{s_i\}_{i \in [3]})$. Then, \mathcal{B}_0 sends $\text{sk}'_f = (\{\text{sk}_{f,i}\}_{i \in [3]}, \pi_{\text{sk}})$ as a reply to \mathcal{A} .

Whenever \mathcal{A} submits a decryption query $(f, \text{ct}' = (\{\text{ct}_i\}_{i \in [2]}, \pi_{\text{ct}}))$, \mathcal{B}_0 generates the functional keys $\text{sk}_{i,f} \leftarrow \text{KeyGen}(\text{mpk}_i, \text{msk}_i, f)$ for $i \in [3]$ and executes $\text{NIZK.Verify}(\text{CRS}, (\text{mpk}_i, \text{ct}_i)_{i \in [3]})$. If $\text{NIZK.Verify}(\text{CRS}, (\text{mpk}_i, \text{ct}_i)_{i \in [2]}) = 1$, \mathcal{B}_0 computes $y_{f,i} := \text{Dec}(\text{mpk}_i, f, \text{sk}_{f,i}, \text{ct}_i)$ for $i \in [3]$ and sends the majority vote, $y \leftarrow \text{MajVal}(y_{f,1}, y_{f,2}, y_{f,3})$, to \mathcal{A} . If the verification outputs 0, \mathcal{B}_0 sends \perp to \mathcal{A} .

This covers the simulation of the game G_β . Finally \mathcal{B}_0 outputs the same bit β' returned by \mathcal{A} . It follows, from the perfect simulation, that the advantage of \mathcal{B}_0 is the same as the advantage of \mathcal{A} . □

As in [52] and in the proof of Theorem 6.6, we prove a corollary that shows that whenever a decryption oracle query is asked and this query contains a valid NIZK proof, then the corresponding ciphertext is explainable under the queried function. This is necessary for the proof of the transition from G_1 to G_2 for the simulation of the decryption oracle.

Claim 6.16. For any PPT adversary \mathcal{A} participating in $G_{1,k}$ for $k \in [3]$, the probability that, during the experiment, \mathcal{A} queries its decryption oracle QDec with a function-ciphertext-pair that is not explainable but has an accepting proof is negligible. Namely, we exhibit a PPT adversary \mathcal{B}_1 , such that

$$\Pr \left[\begin{array}{l} \exists(f, \{\text{ct}'_i\}_{i \in [3]}, \pi') \in Q : (\{\text{ct}'_i\}_{i \in [3]}, \pi') \neq (\{\text{ct}_i\}_{i \in [3]}, \pi), \\ \text{NIZK.Verify}(\text{CRS}, \{\text{ct}'_i\}_{i \in [3]}, \pi') = 1 \text{ and for all } i, j \in [3], \\ i \neq j : \text{Dec}(\text{mpk}_i, \text{sk}_{f,i}, \text{ct}'_i) \neq \text{Dec}(\text{mpk}_j, \text{sk}_{f,j}, \text{ct}'_j) \end{array} \right] \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_1}^{\text{Sim-Sound}}(\lambda),$$

where $\text{sk}_{f,i} \leftarrow \text{KeyGen}(\text{mpk}_i, \text{msk}_i, f)$ for $i \in [3]$, $(\{\text{ct}_i\}_{i \in [3]}, \pi)$ is the reply to the encryption query (x^0, x^1) made by \mathcal{A} . where $(\{\text{ct}_i\}_{i \in [3]}, \pi)$ is the reply to the encryption query (x^0, x^1) made by \mathcal{A} and Q the list containing all the decryption queries $(f, \{\text{ct}'_i\}_{i \in [2]}, \pi')$ asked by \mathcal{A} , knowing the master public key $\text{mpk}' := (\text{CRS}, \{\text{mpk}_i\}_{i \in [2]})$, the reply to its challenge query $(\{\text{ct}_i\}_{i \in [3]}, \pi)$ and by having access to the key generation oracle $\text{KeyGen}'(\text{mpk}', \text{msk}', \cdot)$, during the game.

Proof. We build an adversary \mathcal{B}_1 that simulates $G_{1,k}$ towards \mathcal{A} when interacting with the underlying one-time simulation-soundness experiment.

After the adversary \mathcal{B}_1 has received CRS from the underlying experiment, it generates $(\text{mpk}_i, \text{msk}_i) \leftarrow \text{Setup}(1^\lambda; s_i)$ with $s_i \leftarrow \{0, 1\}^\lambda$ for $i \in [3]$, sets $\text{mpk}' := (\text{CRS}, \{\text{mpk}_i\}_{i \in [3]})$ and sends mpk' to \mathcal{A} . Whenever \mathcal{A} submits a key generation query f , \mathcal{B}_1 generates the functional keys $\text{sk}_{f,i} \leftarrow \text{KeyGen}(\text{mpk}_i, \text{msk}_i, f; r_i)$ for $i \in [3]$ and creates a NIWI proof π_{sk} over the relation R_{set} for the statement-witness pair $y = (\{\text{mpk}_i\}_{i \in [3]}, \{\text{sk}_{f,i}\}_{i \in [3]}, f)$ and $w = (\{\text{msk}_i\}_{i \in [3]}, \{r_i\}_{i \in [3]}, \{s_i\}_{i \in [3]})$. Then, \mathcal{B}_1 sends $\text{sk}'_f = (\{\text{sk}_{f,i}\}_{i \in [3]}, \pi_{\text{sk}})$ as a reply to \mathcal{A} .

For the challenge query (x^0, x^1) asked by \mathcal{A} , \mathcal{B}_1 computes $\text{ct}_i = \text{Enc}(\text{mpk}_1, x^1)$ for $i \leq k$ and $\text{ct}_i = \text{Enc}(\text{mpk}_i, x^0)$ for $i > k$ and asks its experiment for a simulated proof π of the statement $(\text{mpk}_i, \text{ct}_i)_{i \in [3]}$. It sets $\text{ct}' := (\{\text{ct}_i\}_{i \in [3]}, \pi)$ and sends ct' to \mathcal{A} .

Whenever \mathcal{A} outputs a decryption query $(f, \text{ct}' := (\{\text{ct}_i\}_{i \in [3]}, \pi))$, \mathcal{B}_1 verifies the proof. If the output of the verification is 1, \mathcal{B}_1 computes $y_{f,i} = \text{Dec}(\text{mpk}_i, \text{sk}_{f,i}, \text{ct}_i)$ for all $i \in [3]$. Then, \mathcal{B}_1 computes $y \leftarrow \text{MajVal}(y_{f,1}, y_{f,2}, y_{f,3})$ and if $y = \diamond$ then \mathcal{B}_1 sends $(\{\text{ct}_i\}_{i \in [3]}, \pi)$ as a proof forgery to its challenger. Otherwise it sends y to \mathcal{A} . If the verification outputs 0, \mathcal{B}_1 sends \perp to \mathcal{A} . \square

After introducing and proving Claim 6.16, we prove the transition from G_1 to G_2

Lemma 6.17 (Transition from G_1 to G_2). For any PPT adversary \mathcal{A} , there exist PPT adversaries \mathcal{B}_1 and \mathcal{B}_2 , such that

$$|\text{Win}_{\mathcal{A}}^{G_1}(1^\lambda) - \text{Win}_{\mathcal{A}}^{G_2}(1^\lambda)| \leq 3 \cdot \text{Adv}_{\text{NIZK}, \mathcal{B}_1}^{\text{Sim-Sound}}(\lambda) + 3 \cdot \text{Adv}_{\text{FE}, \mathcal{B}_2}^{\text{IND-CPA}}(\lambda) .$$

Proof. To prove that G_1 is indistinguishable from G_2 we apply a hybrid argument over the three slots.

Using the definition of the games in Fig. 25, we can see that

$$|\text{Win}_{\mathcal{A}}^{G_1}(1^\lambda) - \text{Win}_{\mathcal{A}}^{G_2}(1^\lambda)| \leq \sum_{k=1}^3 |\text{Win}_{\mathcal{A}}^{G_{1,k-1}}(1^\lambda) - \text{Win}_{\mathcal{A}}^{G_{1,k}}(1^\lambda)| ,$$

where G_1 corresponds to $G_{1,0}$ and whereas G_2 is identical to game $G_{1,3}$.

Now, we bound the difference between each consecutive pair of games for all $k \in [3]$:

Corollary 6.18. For every $k \in [3]$, there exist PPT adversaries \mathcal{B}_1 and $\mathcal{B}_{2,k}$ such that

$$|\text{Win}_{\mathcal{A}}^{G_{1,k-1}}(1^\lambda) - \text{Win}_{\mathcal{A}}^{G_{1,k}}(1^\lambda)| \leq \text{Adv}_{\text{NIZK}, \mathcal{B}_1}^{\text{Sim-Sound}}(\lambda) + \text{Adv}_{\text{FE}, \mathcal{B}_{2,k}}^{\text{IND-CPA}}(\lambda) .$$

The adversary \mathcal{B} of the statement is then defined as the monolithic adversary that first samples $k \leftarrow [3]$ uniformly at random and then runs the code of $\mathcal{B}_{1,k}$.

Proof. We build an adversary $\mathcal{B}_{2,k}$ that simulates $\mathsf{G}_{1,(k-1+\beta)}$ to \mathcal{A} when interacting with the underlying IND-CPA $_{\beta}^{\text{FE}}$ experiment.

In the beginning of the reduction, $\mathcal{B}_{2,k}$ receives mpk_k from the underlying experiment. It simulates a CRS, i.e. $(\text{CRS}, \tau) \leftarrow \mathcal{S}_1(1^\lambda)$, generates several functional encryption instances $(\text{mpk}_i, \text{msk}_i) \leftarrow \text{Setup}(1^\lambda; s_i)$ with $s_i \in \{0, 1\}^\lambda$ for $i \in [3] \setminus \{k\}$, sets $\text{mpk}' := (\text{CRS}, \{\text{mpk}_i\}_{i \in [3]})$ and sends mpk' to \mathcal{A} . Whenever \mathcal{A} submits a key generation query f , $\mathcal{B}_{2,k}$ forwards this query to its own key generation oracle $\text{KeyGen}(\text{mpk}_k, \text{msk}_k, \cdot)$, to receive $\text{sk}_{f,k}$ as an answer. Then, $\mathcal{B}_{2,k}$ generates $\text{sk}_{f,i} \leftarrow \text{KeyGen}(\text{mpk}_i, \text{msk}_i, f; r_i)$ for $i \in [3] \setminus \{k\}$ by itself and creates a NIWI proof π_{sk} over the relation R_{set} for the statement-witness pair $y = (\{\text{mpk}_i\}_{i \in [3]}, \{\text{sk}_{f,i}\}_{i \in [3]}, f)$ and $w = (\{\text{msk}_i\}_{i \in [3]}, \{r_i\}_{i \in [3]}, \{s_i\}_{i \in [3]})$. Then, $\mathcal{B}_{2,k}$ sends $\text{sk}'_f = (\{\text{sk}_{f,i}\}_{i \in [3]}, \pi_{\text{sk}})$ as a reply to \mathcal{A} .

For the challenge query (x^0, x^1) asked by \mathcal{A} , $\mathcal{B}_{2,k}$ forwards it to its own encryption oracle and receives $\text{ct}_k = \text{Enc}(\text{mpk}_k, x^\beta)$ as an answer. It generates $\text{ct}_i = \text{Enc}(\text{mpk}_i, x^1; u_i)$ with $u_i \leftarrow \{0, 1\}^\lambda$, for $i < k$, and $\text{ct}_i = \text{Enc}(\text{mpk}_i, x^0; u_i)$ with $u_i \leftarrow \{0, 1\}^\lambda$, for $i > k$, by itself, simulates a valid proof π of the relation $R_{\text{set}}^{\text{CCA}}$ using the statement $y = (\text{mpk}_i, \text{ct}_i)_{i \in [3]}$, i.e. $\pi_{\text{ct}} \leftarrow \mathcal{S}_2(\text{CRS}, \tau, y)$ and sends $\text{ct}' := (\{\text{ct}_i\}_{i \in [3]}, \pi_{\text{ct}})$ to \mathcal{A} .

Whenever \mathcal{A} asks a decryption query $(f, \text{ct}' := (\{\text{ct}_i\}_{i \in [3]}, \pi_{\text{ct}}))$, $\mathcal{B}_{2,k}$ first verifies the proof π_{ct} , i.e. it executes $\text{NIZK.Verify}(\text{CRS}, (\text{mpk}_i, \text{ct}_i)_{i \in [2]}, \pi_{\text{ct}})$. If the verification outputs 1, $\mathcal{B}_{2,k}$ generates $\text{sk}_{f,i} \leftarrow \text{KeyGen}(\text{mpk}_i, \text{msk}_i, f)$ for $i \in [3] \setminus \{k\}$, computes $y_{f,i} \leftarrow \text{Dec}(\text{mpk}_i, f, \text{sk}_{f,i}, \text{ct}_i)$ for $i \in [3] \setminus \{k\}$ and $y \leftarrow \text{MajVal}(\{y_{f,i}\}_{i \in [3] \setminus \{k\}})$, to \mathcal{A} . Since Claim 6.16 shows that for all ciphertext queries made by \mathcal{A} that have a valid proof (except with negligible probability), it holds that $\text{Dec}(\text{mpk}_1, f, \text{sk}_{f,1}, \text{ct}_1) = \text{Dec}(\text{mpk}_2, f, \text{sk}_{f,2}, \text{ct}_2) = \text{Dec}(\text{mpk}_3, f, \text{sk}_{f,3}, \text{ct}_3)$. Therefore, it is sufficient to generate the decryptions $y_{f,1}$ and $y_{f,2}$ and use them as the decryption output and a reply for \mathcal{A} . If the verification outputs 0, \mathcal{B}_2 sends \perp to \mathcal{A} .

This covers the simulation of the game $\mathsf{G}_{1,(k-1+\beta)}$. Finally \mathcal{B}_2 outputs the same bit β' returned by \mathcal{A} . Together with the analysis of adversary \mathcal{B}_1 , this yields the advantage mentioned in the lemma. \square

Besides the security lifting, the compiler also achieves the desired setup consistency property. This can be shown as in the proof for Theorem 6.11 .

Theorem 6.19. *Let $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be a functional encryption scheme, $\text{NIWI} = (\text{NIWI.Prove}, \text{NIWI.Verify})$ a NIWI proof system for R_{set} (Fig. 21) and $\text{NIZK} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ a NIZK proof system for $R_{\text{set}}^{\text{CCA}}$ (Fig. 24), then the construction $\text{FE}' = (\text{Setup}', \text{KeyGen}', \text{Enc}', \text{Dec}')$ defined in Fig. 23 is setup consistent. Namely, for any PPT adversary \mathcal{A} , exists a PPT adversary \mathcal{B} such that*

$$|\Pr[\text{set-CONS}^{\text{FE}'}(1^\lambda, \mathcal{A}) = 1] \leq \text{Adv}_{\text{NIWI}, \mathcal{B}}^{\text{Sound}}(\lambda).$$

Proof (Sketch). The proof proceeds following the same reasoning as the proof of Theorem 6.11. The reason is that the introduction of the additional proof of the ciphertext must always yield 1 (as otherwise, the ciphertext will not be considered by set-CONS) and only those ciphertexts are considered by Dec' , as defined in the compiler. Hence, we can invoke the same analysis, based on the invariant that for each ciphertext $\text{ct}'_i = (\text{ct}_i, \pi_i)$ ($i = 1, 2$), π_i is valid therefore we can perform the identical case distinctions as in the proof of Theorem 6.11 based on the ciphertexts ct_i . \square

6.3 Strong Input Consistency

In this section, we show that the verifiability property introduced in [12], can be understood as providing strong input consistency. Due to this modular reduction, we also directly inherit their compiler, and in general any compiler that achieves verifiable functional encryption.

We recall the syntax of VFE in Definition 2.16. In a nutshell, VFE extends standard FE by two additional algorithms:

- VerifyCT(mpk, ct):** A predicate on ciphertexts (w.r.t. the public key) that decides whether ct is valid.
- VerifySK(mpk, f, sk):** A predicate on pairs (sk, f) (w.r.t. the public key) that decides whether the pair is a valid key-function pair.

The *verifiability property* of [12] restated in Definition 2.17 of the supplemental material requires that whenever $\text{VerifyCT}(\text{mpk}, \text{ct}) = \text{VerifySK}(\text{mpk}, f, \text{sk}) = 1$, we have $\Pr[\text{Dec}(\text{mpk}, f, \text{sk}, \text{ct}) = f(x)] = 1$ (where the implication must hold over all possible values of the involved arguments).

We define a simple compiler, defined Fig. 26, in that makes use of these two verification procedures to achieve strong input consistency. Informally, the first algorithm is used as a ciphertext verification check (and we return \perp if the check fails) and the second function is used to verify key-function pairs (and return \diamond if the check fails). Note that the transformation clearly preserves the confidentiality notion of the underlying VFE scheme.

<p><u>Setup'(1λ) :</u> $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ Return $(\text{mpk}', \text{msk}') = (\text{mpk}, \text{msk})$</p> <p><u>KeyGen'(mpk', msk', f) :</u> Parse $\text{mpk}' = \text{mpk}, \text{msk}' = \text{msk}$ $\text{sk}_f = \text{KeyGen}(\text{mpk}, \text{msk}, f)$ Return $\text{sk}'_f = \text{sk}_f$</p> <p><u>Enc'(mpk', x) :</u> Parse $\text{mpk}' = \text{mpk}, \text{ct} = \text{Enc}(\text{mpk}, x)$ Return $\text{ct}' = \text{ct}$</p>	<p><u>Dec'(mpk', f, sk'_f, ct') :</u> Parse $\text{mpk}' = \text{mpk}, \text{sk}'_f = \text{sk}_f, \text{ct}' = \text{ct}$ If $\text{VerifyCT}(\text{mpk}, \text{ct}) \neq 1$ Return \perp If $\text{VerifySK}(\text{mpk}, f, \text{sk}_f) \neq 1$ Return \diamond $y := \text{Dec}(\text{mpk}, f, \text{sk}_f, \text{ct})$ Return y</p>
---	--

Fig. 26: Strong input consistency compiler

Theorem 6.20. *Let VFE = (Setup, KeyGen, Enc, Dec, VerifyCT, VerifySK) be a verifiable functional encryption scheme then the construction FE' = (Setup', KeyGen', Enc', Dec') defined in Figure 26 is strongly input consistent. Namely, for any PPT adversary \mathcal{A} , it holds that:*

$$|\Pr[\text{st-in-CONS}^{\text{FE}'}(1^\lambda, \mathcal{A}, \text{Extract}) = 1]| = 0 ,$$

where **Extract** is the algorithm that always outputs unknown.

Proof. To prove the strong input consistency of the scheme FE', we rely on the verifiability of the VFE scheme. In more detail, we construct an adversary \mathcal{B} that violates verifiability with the probability with which \mathcal{A} wins the strong input consistency experiment st-in-CONS.

In the beginning of the reduction, \mathcal{B} receives a master public key mpk , two ciphertexts ct_1, ct_2 , and a tuple of secret keys with the corresponding functions $\{(\text{sk}_j, f_j)\}_{j \in [n]}$ from \mathcal{A} . In the next step, \mathcal{B} computes $y_{j,i} := \text{Dec}'(\text{mpk}', f_j, \text{sk}_j, \text{ct}_i)$ for all $j \in [n], i \in \{1, 2\}$ and defines the set F as all the functional keys that do not output \diamond , i.e. $F := \{(\text{sk}_j, f_j)\}_{j \in [n] \wedge (y_{j,1} \neq \diamond \vee y_{j,2} \neq \diamond)}$. Let E denote the event that the intersection $\bigcap_{j \in [n], (\cdot, f_j) \in F} f_j^{-1}(y_{j,i})$ is equal to \emptyset . We show that the occurrence of E contradicts the verifiability notion.

For the sake of contradiction, we assume that verifiability holds as stated. Now, we analyze the different scenarios in which the intersection is empty and show that these cannot occur. For this purpose, we define the events E_1, E_2, E_3 :

Event E_1 denotes the case that if $y_{j,i} \neq \perp$ for a single $j \in [n]$ and a fixed $i \in \{0, 1\}$, then $y_{j,i} \neq \perp$ for all $j \in [n]$ and the same fixed $i \in \{0, 1\}$. Event E_2 denotes the case that if $y_{j,i} \neq \diamond$ for a single $i \in \{0, 1\}$ and a fixed $j \in [n]$, then $y_{j,i} \neq \diamond$ for both $i \in \{0, 1\}$ and the same fixed $j \in [n]$. The final event, event E_3 , denotes the case that if $y_{j,i} \notin \{\perp, \diamond\}$, then $y_{j,i} = f_j(x_i)$.

We start by analyzing event E_1 . Let one of the $y_{j,i} = \perp$ for a single $j \in [n]$ and a fixed $i \in \{1, 2\}$. Now, if **VerifyCT** was satisfied by ct_i , then by the verifiability condition, \perp is a valid output and thus by definition of

the special symbol \perp , the only “preimage” explaining the output is \perp and all the decryptions of ct_i under different functional keys would lead the same output. Therefore, under the above assumption, all other $y_{j,i}$ must yield \perp .

In the other case, the `VerifyCT` algorithm outputs 0 for the ciphertext ct_i (for a fixed i). However, then the `VerifyCT` algorithm also outputs 0 on every other decrypt request since it is deterministic and does not depend on any functional key $\text{sk}_{j'}$, $j' \in [n] \setminus \{j\}$. This leads to the fact that $y_{j,i} = \perp$ for all $j \in [n]$ and a fixed $i \in \{1, 2\}$. In this case, the intersection contains the \perp symbol.

For event E_2 , let one of the $y_{j,i} = \diamond$ for a single $i \in \{1, 2\}$ and a fixed $j \in [n]$. If this case occurs, then the `VerifySK` algorithm must have output 0 for the functional key sk_j for a fixed j as otherwise, since by definition of the special symbol \diamond , there is no “preimage” explaining the output (and thus the verifiability property violated). As before, the `VerifySK` algorithm also outputs 0 if it gets queried using another ciphertext $\text{ct}_{i'}$ but the same functional key sk_j . This yields that $y_{j,i'} = \diamond$ for $i' \neq i$ and directly deletes the key sk_j from the list F (due to the definition of F).

Finally, we analyze event E_3 . Let $y_{j,i} \notin \{\perp, \diamond\}$, then both of the verify algorithms, `VerifyCT` and `VerifySK`, output 1 in the decryption procedure. This ensures, together with the verifiability property, i.e. $\Pr[\text{Dec}(\text{mpk}, f_j, \text{sk}, \text{ct}_i) = f_j(x_i)] = 1$, where f_j is the function associated with sk_j and x_i the plaintext associated with ct_i , that $y_{j,i} = f_j(x_i)$.

Taking into account the analysis of event E_1 and E_2 , the decryption $y_{j',i}$, corresponding to a different functional key $\text{sk}_{j'}$, is unequal to \diamond (due to event E_4 and the definition of F) and unequal to \perp (due to event E_1 and the fact that $y_{j,i}$ is a valid decryption). By doing the same analysis for $y_{j',i}$ as for $y_{j,i}$ as in event E_2 , we obtain that $y_{j',i} = f_{j'}(x_i)$. We can do the same analysis for all the remaining $y_{j'',i}$ with $j'' \in [n] \setminus \{j, j'\}$ and therefore it follows that $\bigcap_{j \in [n], (\cdot, f_j) \in F} f_j^{-1}(y_{j,i}) \neq \emptyset$. The same analysis also needs to be done for the second ciphertext $\text{ct}_{i'}$ with $i' \neq i$.

This shows that event E cannot occur and therefore that the proposed construction achieves strong input consistency. \square

By showing the relation above, our treatment nicely includes the verifiability property and hence the analysis in the next section also gives a UC interpretation of the construction (achieved by both strong input consistency and verifiability).

In comparison, our strong input consistency notion is more positioned as a property that an attacker tries to break, rather than a general requirement of a scheme that holds for all arguments and follows directly as a strengthening of input consistency.

While the verifiability property is also necessary for our concrete compiler (as long as `Dec` is deterministic), strong input consistency does not achieve the verifiability property on its own. The reason for this is that the definition of strong input consistency does not change the FE syntax, verification checks are inherently done by `Dec` with access to at least one secret key. Thus, no guaranteed verifiability algorithm of the form `VerifyCT(mpk, ct)` (or `VerifySK(mpk, f, sk_f)`) can be directly deduced from a generic FE scheme that is strongly input-consistent according to our notion. Hence, our notion puts forth a seemingly weaker form of (implicit) secret-key verifiability.

Also, when leaving the standard model (e.g. switching to the random-oracle model), Definition 2.17 could technically be violated even though *finding* these values by an efficient adversary \mathcal{A} —which the strong input consistency definition asks for—might be infeasible. We further observe as an interesting open problem, whether strong input consistent (and also verifiable) FE schemes exist that satisfy CCA or CFE security (e.g., in the random-oracle model).

7 UC Consistency for Functional Encryption

Thanks to the foundational work of Matt and Maurer [53], we can accurately characterize the goal that functional encryption tries to achieve in a constructive sense: the goal is to realize a repository, where an input provider can store values x —and receive a unique handle h to access it—and where another entity can, using the handle h , obtain the values $f(x)$ if and only if the function f was explicitly granted for evaluation

by a third-party (the setup manager). This model puts forth a form of access control: even a malicious receiver (or a coalition of malicious receivers) cannot obtain more information than what $f(x)$ reveals. The goal of plain functional encryption is to realize this strong repository from a public one, where the receivers have read access to all values stored in the repository. While the treatment in [53] focuses on confidentiality (in the sense that only the receiver is corrupted), we show how to extend the treatment to all other cases, including malicious input providers and/or malicious setup generators. We thereby extend their functionality to precisely capture the realistic composable notion of consistency in functional encryption. We then prove that the security game we put forward in the previous sections are adequate characterization of this ideal functionality. Furthermore, we directly model the case of multiple receivers.

In the following, we let \mathcal{F} denote the concrete functionality class of an FE scheme (dropping the index λ). We denote by f_0 the distinguished leakage function and let $\mathcal{F}^+ := \mathcal{F} \setminus \{f_0\}$. We denote by \mathcal{X} the input domain and by \mathcal{Y} the range of the functions. Further, let \mathbf{Q} be a monotone predicate on function sets of \mathcal{F} . We further assume three distinct identifiers $\mathbf{A}, \mathbf{B}, \mathbf{C}$. Furthermore, if \mathbf{x} is an identifier, \mathbf{x}_i denotes the unique identifier derived from \mathbf{x} that includes a (prefix-free) encoding of the the number i .

The Ideal Functionality. Our ideal functionality $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathbf{Q})}^{\mathbf{A}, \mathbf{B}, \mathbf{C}, t}$ is a natural relaxation of the one provided in [53] and is formally described below. The functionality is defined to interact with three types of parties or roles, denoted by \mathbf{A} for the input provider, \mathbf{B} for the receiver or decryptor, and \mathbf{C} for the third party that manages setup and key distribution. The ideal functionality provides perfect confidentiality by only allowing the corrupted receivers to learn the function values $f(x)$ for inputs x of the input provider and functions f assigned to the corrupted receivers.

Turning towards UC consistency, we see that it is a property that protects an *honest* receiver against malicious input providers and/or setup generator: an honest receiver, when requested to provide the function value corresponds to an entry in the repository, will always output values that are consistent with any function he was assigned, i.e., he will never encounter a situation, in which a set of output values would indicate that the underlying value in the repository can actually not exist. We separate two important aspects here:

- Consistency is the property that the input provider, at any time, is committed per handle h to a set of possible values $\mathcal{M} \subseteq \mathcal{X}$ that are consistent with any output y generated by an honest receiver when requesting the output of the function value for an assigned function f for handle h .
- Commitment to exactly one value per handle is captured via predicate $\mathbf{Q}()$: if the set of functions R assigned to the receiver satisfies $\mathbf{Q}(R)$, then the functionality enforces that to each handle h only one value $x \in \mathcal{X}$ exists and must be fixed even if Alice is corrupted (if Alice is honest, the value is the value provided by Alice). Once $\mathbf{Q}(R) = 1$, the repository has an identical behavior to the one defined in [53]. Note that the predicate allows us to have a more realistic notion (because the task of a simulator is easier since it does not need to extract the value x immediately when seeing a ciphertext).

Overview of the technical description. If \mathbf{A} is corrupted, we allow it to not only provide inputs in the domain, but also to specify that the input is not fixed ($x = \text{unknown}$), in which case the functionality associates with a handle h a set of allowed values (initially the entire domain). However, these inputs are only allowed if the set of functions assigned to the receivers \mathbf{B}_i does not satisfy predicate $\mathbf{Q}(R)$, in which case the functionality demands that the value assigned to handle h must be a singleton set. The functions assigned to the receivers are steered by party \mathbf{C} . If this party is corrupted, the capability the setup manager has is to stop the ideal functionality from accepting Alice’s inputs (this is unavoidable as Alice waits for the public parameters of the system anyway). For an honest party \mathbf{B}_i , the ideal functionality implements the following guarantees: whenever requested to obtain an output value corresponding to a function f and handle h , $f(x)$ is returned, if the handle h is associated to a unique element x . Otherwise, the functionality lets the adversary choose the return value y upon this request but enforces that the set of values assigned to handle h is reduced accordingly to stay consistent with this output. For the sake of generality—and in order to make sure that the behavior of the functionality can be captured as an efficient program as required in UC—we assume an efficiently computable map $\text{preMap}()$ at this step satisfying $\text{preMap}(\mathcal{M}, f, y) = \{x' \in \mathcal{M} \mid f(x') = y\}$ to formulate the ideal consistency guarantees. Finally, if the predicate $\mathbf{Q}(R)$ becomes satisfied at any point

during the execution, then the functionality enforces that the adversary commits to a unique value x at this point. Additionally, as in [53], we assume a function `getHandle` that returns unique identifiers (but abstract our treatment from a concrete instantiation).

We note that the above consistency checks are only enforced if party B_i is honest. If only receivers are dishonest (i.e., the input provider and the setup manager are honest) then the functionality outputs $f(x)$ for the unique x input by honest party A if and only if f was either assigned by the setup manager C or if the leakage function $f = f_0$ is queried (which is exactly the behavior implemented in [53]). Last but not least, note that we do not provide ideal confidentiality guarantees when either the input provider or the setup manager are dishonest (e.g., a malicious setup generator could produce arbitrary secret keys and collude with the corrupted receivers).

Functionality $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\text{A}, \text{B}, \text{C}, t}$

The functionality is parameterized by the function class \mathcal{F} , a predicate $\mathcal{Q}(\cdot) : 2^{\mathcal{F}} \rightarrow \{0, 1\}$, and by three distinct party identities (of the dummy parties) interacting with the functionality denoted by $\mathcal{P} := \{\text{A}, \text{B}, \text{C}\}$ (that identify a particular role), and the number t of decryptors/receivers.

Setup. Upon receiving input $(\text{SETUP}, \text{sid})$ via dummy party C (or from the adversary on behalf of corrupted C), set $\text{setup} \leftarrow \text{true}$, $R_i \leftarrow \emptyset$, for each $i \in [t]$. Ignore the request if the party-id does not correspond to C. Output $(\text{SETUP}, \text{sid})$ to the adversary to indicate that setup is completed.

Input. Upon receiving input $(\text{WRITE}, \text{sid}, x)$ via dummy party A (or from the adversary on behalf of corrupted A), and if $\text{setup} = \text{true}$, do the following:

- If pid A is honest then verify that $x \in X$ (ignore request otherwise). If party C is honest, then compute handle $h \leftarrow \text{getHandle}$ and store $M[h] \leftarrow (x, \{x\})$. Return $(\text{WRITTEN}, \text{sid}, h)$ to the calling party. If party C is corrupted, then provide public delayed-output to the adversary and do the previous actions only upon receiving ACK for this operation.^a
- If pid A is marked as corrupted and $\mathcal{Q}(\cup_{i=1}^t R_i)$ holds, verify that $x \in X \cup \{\perp\}$ (and ignore the request otherwise). Choose $h \leftarrow \text{getHandle}$ and store $M[h] \leftarrow (x, \{x\})$. Output $(\text{WRITTEN}, \text{sid}, h)$ to the adversary.
- If pid A is marked as corrupted and $\neg \mathcal{Q}(\cup_{i=1}^t R_i)$, just verify that $x \in X \cup \{\perp, \text{unknown}\}$ (and ignore the request otherwise) and perform the same actions as for the previous case.

Access Management. Upon receiving input $(\text{ASSIGN}, \text{sid}, f, i)$ via dummy party C (or from the adversary on behalf of C), do the following: if $f \in \mathcal{F}^+$, then update $R_i \leftarrow R_i \cup \{f\}$ and output $(\text{ASSIGNED}, \text{sid}, f, i)$ to the adversary.

Output. Upon receiving $(\text{READ}, \text{sid}, h, f)$ from some caller via dummy party B_i (or from the adversary on behalf of corrupted B), first parse $M[h]$ as (x, \mathcal{M}) . In case $M[h] = \perp$, return `noData`.

- If B_i is honest do:
 1. If $f \notin R_i$ then give up activation. Otherwise, if $x \in X$ and $f \in R_i$, then return $(\text{READ}, \text{sid}, f(x))$ to the caller; else if $x = \perp$ then return $(\text{READ}, \text{sid}, \perp)$ to the calling party. (***)
 2. Otherwise, if $x = \text{unknown}$, output $(\text{READ}, \text{sid}, h, f)$ to the adversary. Upon receiving $(\text{READ}, \text{sid}, h, f, (x, y))$ from the adversary, do the following:
 - (a) If $x \in \mathcal{M}$, set $M[h] \leftarrow (x, \{x\})$. Output $(\text{READ}, \text{sid}, f(x))$ to the calling party.
 - (b) Else, if $\mathcal{Q}(\cup_{i=1}^t R_i)$ (but $x \notin \mathcal{M}$) then pick x' at random from \mathcal{M} and store $M[h] \leftarrow (x', \{x'\})$. Output $(\text{READ}, \text{sid}, f(x'))$ the calling party. (*)
 - (c) Else (i.e. $\neg \mathcal{Q}(\cup_{i=1}^t R_i)$ and $x \notin \mathcal{M}$) compute $\mathcal{M}_{\text{new}} \leftarrow \text{preMap}(\mathcal{M}, f, y)$.
 - i. If $\mathcal{M}_{\text{new}} = \emptyset$ then pick x' at random from \mathcal{M} and store $M[h] \leftarrow (x', \{x'\})$. Output $(\text{READ}, \text{sid}, f(x'))$ the calling party. (**)

- ii. Otherwise, update $M[h] \leftarrow (\text{unknown}, \mathcal{M}_{\text{new}})$ and output $(\text{READ}, \text{sid}, y)$ to the calling party.
- If B_i is marked as corrupted but none of A or C , then do:
 1. If $f \in R_i$ then return $(\text{READ}, \text{sid}, f(x))$ (for the x guaranteed to exist since the input provider is honest) and if $f = f_0$ then return $(\text{READ}, \text{sid}, f_0(x))$ to the adversary. Otherwise, give up activation.
- If B_i is corrupted alongside A or C , then output $(\text{READ}, \text{sid}, h, f)$ to the adversary and upon receiving $(\text{READ}, \text{sid}, h, f, (x, y))$ from the adversary output $(\text{READ}, \text{sid}, y)$ to the calling party.

Additional adversarial interaction.

- On top of the standard pid-wise corruption mechanism of UC, the following additional capability is given to the adversary: If and only if some B_i and at least one more party among $\{A, C\}$ is corrupted, then the adversary is allowed to query $(\text{REVEAL}, \text{sid}, h)$ upon which $M[h]$ is revealed to the adversary.

^a See also Remark 7.1 for a strengthened version of this case.

Remark 7.1 (Ideal Universal Encryption). It is interesting to see how the additional *universal encryption property*, as discussed in Definition 3.5, would strengthen the ideal world: It would reduce the power of the adversary when corrupting the setup manager C : instead of the capability to decide per message whether a write request by the input provider succeeds (which is formalized above), the adversary could only decide *once and for all* at the point the first message is written whether the repository is “writable” or not. Hence, we can get rid of the delayed output which might be extremely useful for certain applications (and motivated our game-based property of universal encryption). We prove here the weaker version, as most of our constructions do not have the universal encryption property (as discussed in the respective sections).

The FE Protocol. We define the protocol $\pi_{\text{FE}}^{\text{A,B,C},t}$ for parties A , B_i and C , where party A acts as an input provider, parties B_i act as the receivers, and party C acts as the setup manager. The distribution of the public key to the input provider and the receivers is done via an authentic broadcast channel between C and A and between C and the receivers $\{B_i\}_{i \in [t]}$. Note that we do not require broadcast from C to every participant, but only to the set of parties having the same role, which is the minimal assumption we have to make. (To see this, note that otherwise, there could be two parties with the same role that operate with public parameters belonging to different schemes—among which clearly no consistency has to exist: for example when one scheme deems all legitimate encryptions w.r.t. to the other public key invalid under its own public key.) In addition, the protocol requires a point-to-point secret channel between C and each of the receivers B_i . Finally, we assume a basic storage repository, where the input provider (and only the input provider) can store messages of its choice (and only the receivers B_i can access them). Note that in UC, these hybrid functionalities are defined and invoked by the protocol. Hence, if a scheme would require the random oracle model, $\pi_{\text{FE}}^{\text{A,B,C},t}$ would additionally invoke a random-oracle functionality (which is needed to achieve CFE security for example). The channel functionalities and the basic real-world repository are given in Appendix C.1 for completeness. In a nutshell, the protocol works as follows: Party C generates the public keys (and sends them to the other parties) and assigns functions to parties B_i by sending the functional keys. Party A does provide the input to the real-world repository by encrypting the input $x \in \mathcal{X}$, and storing valid ciphertexts in the repository. Using the obtained handle h , the ciphertext can be accessed by some party B_i and decrypted using a (valid) sk corresponding to an assigned function f and the result is provided as output. The protocol is specified below.

UC Realization. We provide a detailed security analysis with respect to the different corruption sets possible in the system and conclude that each of our consistency games captures exactly what we intended. The theorem therefore also gives guarantees for a scheme that does only achieve a subset of the properties

Protocol $\pi_{\text{FE}}^{A,B,C,t}$

- Upon each invocation, protocol $\pi_{\text{FE}}^{A,B,C,t}$ first verifies that this ITI's party identifier matches $\text{pid} \in \{A, B_i, C\}$ ($i \in [t]$) and rejects the message otherwise. This means that the ITI running the protocol must have the extended identity $\text{eid}_{\text{pid}} = (\pi_{\text{FE}}^{A,B,C,t}, \text{sid} || \text{pid})$ for some sid and $\text{pid} \in \{A, B_i, C\}$ (for some $i \in [t]$).
- Depending on the encoded pid , match the input to the following commands:
 - pid = C:** The behaviour of the manager is as follows:
 - On input (SETUP, sid) execute $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}()$, store the pair internally and send mpk via $\text{Func}_{\text{auth}}^{\text{eid}_{\text{C}}, \text{eid}_{\text{A}}}$ to the input provider and via $\text{Func}_{\text{auth}}^{\text{eid}_{\text{C}}, \text{eid}_{\text{B}_1}, \dots, \text{eid}_{\text{B}_t}}$ to the decryptor.
 - On input (ASSIGN, sid, f, i), ensure that $f \in F$ and otherwise ignore the input. Execute $\text{sk}_f = \text{KeyGen}(\text{mpk}, \text{msk}, f)$, send sk_f via $\text{Func}_{\text{sec}}^{\text{eid}_{\text{C}}, \text{eid}_{\text{B}_i}}$ to the decryptor.
 - pid = A:** The behaviour of the input provider is as follows:
 - On input (WRITE, sid, x), ensure that $x \in \mathcal{X}$ and that an mpk has been received (otherwise ignore the input). Then, execute $\text{ct} \leftarrow \text{Enc}(\text{mpk}, x)$ and if $\text{ct} \neq \text{err}$ output (WRITE, sid, ct) to $\text{Func}_{\text{basic-rep}, \text{C}}^{A, B, t}$ and return the obtained handle h from the basic repository back to the caller by returning (WRITTEN, sid, h). (Give up activation if an error occurs).
 - On receiving the master public key mpk from $\text{Func}_{\text{auth}}^{\text{eid}_{\text{C}}, \text{eid}_{\text{A}}}$ and if this is the first time the key is delivered, store it internally for future reference. Ignore any future message from the channel.
 - pid = B_i:** The behaviour of the decryptor is as follows:
 - On input (READ, sid, h, f) output (READ, sid, h) to $\text{Func}_{\text{basic-rep}, \text{C}}^{A, B, t}$ to obtain a ciphertext ct . If no ciphertext is received or a pair (sk_f, f) is not recorded, then give up activation. Next, execute $y \leftarrow \text{Dec}(\text{mpk}, f, \text{sk}_f, \text{ct})$ for each candidate pair (sk_f, f) recorded and delete the pair if $y = \diamond$ is obtained. Finally, give up the activation if all values returned \diamond . Otherwise, return the first $y \in \mathcal{Y} \cup \{\perp\}$ obtained by decrypting and output (READ, sid, y) to the caller.
 - On receiving the master public key mpk from $\text{Func}_{\text{auth}}^{\text{eid}_{\text{C}}, \text{eid}_{\text{B}_1}, \dots, \text{eid}_{\text{B}_t}}$ and if this is the first time the master public key is delivered, store it internally for future reference.
 - On receiving a pair (sk, f) from $\text{Func}_{\text{sec}}^{\text{eid}_{\text{C}}, \text{eid}_{\text{B}_i}}$ and if $f \in F$, then store the pair (sk, f) in the list of received functional keys.
- Ignore the input if no case applies.

(such as CFE and setup or input consistency): in this case, the scheme can only be safely used in contexts, where certain people are trusted.¹³

Theorem 7.2. $\text{FE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ be a functional encryption scheme for functionality class \mathcal{F} , let Q be a monotone predicate $2^{\mathcal{F}} \rightarrow \{0, 1\}$, and let A, B, C be three identifiers. Protocol $\pi_{\text{FE}}^{A,B,C,t}$ UC-realizes $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, Q)}^{A,B,C,t}$ (under static corruption) under the following conditions:

- If party A is corrupted, and C is honest (and potentially a subset of receivers is corrupted), then in-CONS w.r.t. a Q -extractable algorithm Extract is a sufficient requirement on FE such that $\pi_{\text{FE}}^{A,B,C,t}$ realizes $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, Q)}^{A,B,C,t}$.
- If party C is corrupted and party A is honest (and possibly a subset of receivers are corrupted), then set-CONS is a sufficient requirement on FE such that $\pi_{\text{FE}}^{A,B,C,t}$ realizes $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, Q)}^{A,B,C,t}$.
- If parties A and C are corrupted (and possibly a subset of receivers are corrupted), then st-in-CONS w.r.t. a Q -extractable algorithm Extract is a sufficient requirement on FE such that $\pi_{\text{FE}}^{A,B,C,t}$ realizes $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, Q)}^{A,B,C,t}$.
- If both A and C are honest, and only a subset of receivers is corrupted, then CFE security is a sufficient requirement on FE such that $\pi_{\text{FE}}^{A,B,C,t}$ realizes $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, Q)}^{A,B,C,t}$.

¹³ Technically, this could even be enforced by the UC model in that certain parties are marked incorruptible and the corresponding protocol would ignore corruption requests.

Additionally, in-CONS, set-CONS, and st-in-CONS are the respectively necessary requirements on the scheme FE in order for $\pi_{\text{FE}}^{A,B,C,t}$ to realize the specified security guarantees by $\text{Func}_{\text{Rep},(\mathcal{F}^+,f_0,Q)}^{A,B,C,t}$ w.r.t. a given corruption set in the above listing.

Note that the second part of the theorem justifies our game-based notions for consistency and matches our motivation. We refer to Appendix C for the proof.

Acknowledgments. This work was supported in part by the European Union’s Horizon 2020 Research and Innovation Programme under grant agreement 780108 (FENTEC). C.B.: Work done while author was at The University of Edinburgh, Scotland, UK.

References

1. M. Abdalla, M. Bellare, and G. Neven. Robust encryption. In D. Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 480–497. Springer, Heidelberg, Feb. 2010. (Pages 2 and 3.)
2. M. Abdalla, F. Benhamouda, and R. Gay. From single-input to multi-client inner-product functional encryption. In S. D. Galbraith and S. Moriai, editors, *ASIACRYPT 2019, Part III*, volume 11923 of *LNCS*, pages 552–582. Springer, Heidelberg, Dec. 2019. (Page 58.)
3. M. Abdalla, F. Benhamouda, M. Kohlweiss, and H. Waldner. Decentralizing inner-product functional encryption. In D. Lin and K. Sako, editors, *PKC 2019, Part II*, volume 11443 of *LNCS*, pages 128–157. Springer, Heidelberg, Apr. 2019. (Page 58.)
4. M. Abdalla, D. Catalano, D. Fiore, R. Gay, and B. Ursu. Multi-input functional encryption for inner products: Function-hiding realizations and constructions without pairings. In H. Shacham and A. Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 597–627. Springer, Heidelberg, Aug. 2018. (Page 8.)
5. M. Abdalla, R. Gay, M. Raykova, and H. Wee. Multi-input inner-product functional encryption from pairings. In J. Coron and J. B. Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 601–626. Springer, Heidelberg, Apr. / May 2017. (Page 23.)
6. S. Agrawal, S. Gorbunov, V. Vaikuntanathan, and H. Wee. Functional encryption: New perspectives and lower bounds. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 500–518. Springer, Heidelberg, Aug. 2013. (Page 1.)
7. S. Agrawal, V. Koppula, and B. Waters. Impossibility of simulation secure functional encryption even with random oracles. In A. Beimel and S. Dziembowski, editors, *TCC 2018, Part I*, volume 11239 of *LNCS*, pages 659–688. Springer, Heidelberg, Nov. 2018. (Page 1.)
8. S. Agrawal, B. Libert, and D. Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 333–362. Springer, Heidelberg, Aug. 2016. (Pages 3, 23, 24, and 37.)
9. P. Ananth, D. Boneh, S. Garg, A. Sahai, and M. Zhandry. Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689, 2013. <http://eprint.iacr.org/2013/689>. (Page 58.)
10. P. Ananth and A. Jain. Indistinguishability obfuscation from compact functional encryption. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 308–326. Springer, Heidelberg, Aug. 2015. (Page 58.)
11. B. Auerbach, M. Bellare, and E. Kiltz. Public-key encryption resistant to parameter subversion and its realization from efficiently-embeddable groups. In M. Abdalla and R. Dahab, editors, *PKC 2018, Part I*, volume 10769 of *LNCS*, pages 348–377. Springer, Heidelberg, Mar. 2018. (Page 15.)
12. S. Badrinarayanan, V. Goyal, A. Jain, and A. Sahai. Verifiable functional encryption. In J. H. Cheon and T. Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 557–587. Springer, Heidelberg, Dec. 2016. (Pages 2, 3, 4, 5, 10, 17, 46, 47, and 58.)
13. S. Badrinarayanan, D. Gupta, A. Jain, and A. Sahai. Multi-input functional encryption for unbounded arity functions. In T. Iwata and J. H. Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 27–51. Springer, Heidelberg, Nov. / Dec. 2015. (Page 58.)
14. B. Barak, S. J. Ong, and S. P. Vadhan. Derandomization in cryptography. *SIAM J. Comput.*, 37(2):380–400, 2007. (Pages 3, 8, 9, and 10.)
15. M. Barbosa and P. Farshim. On the semantic security of functional encryption schemes. In K. Kurosawa and G. Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 143–161. Springer, Heidelberg, Feb. / Mar. 2013. (Page 1.)

16. M. Bellare, G. Fuchsbauer, and A. Scafuro. NIZKs with an untrusted CRS: Security in the face of parameter subversion. In J. H. Cheon and T. Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 777–804. Springer, Heidelberg, Dec. 2016. (Page 15.)
17. M. Bellare and A. O’Neill. Semantically-secure functional encryption: Possibility results, impossibility results and the quest for a general definition. In M. Abdalla, C. Nita-Rotaru, and R. Dahab, editors, *CANS 13*, volume 8257 of *LNCS*, pages 218–234. Springer, Heidelberg, Nov. 2013. (Page 1.)
18. M. Ben-Or, O. Goldreich, S. Goldwasser, J. Hästad, J. Kilian, S. Micali, and P. Rogaway. Everything provable is provable in zero-knowledge. In S. Goldwasser, editor, *CRYPTO’88*, volume 403 of *LNCS*, pages 37–56. Springer, Heidelberg, Aug. 1990. (Page 8.)
19. F. Benhamouda, F. Bourse, and H. Lipmaa. CCA-secure inner-product functional encryption from projective hash functions. In S. Fehr, editor, *PKC 2017, Part II*, volume 10175 of *LNCS*, pages 36–66. Springer, Heidelberg, Mar. 2017. (Pages 1 and 6.)
20. D. J. Bernstein and T. Lange. Computing small discrete logarithms faster. In S. D. Galbraith and M. Nandi, editors, *INDOCRYPT 2012*, volume 7668 of *LNCS*, pages 317–338. Springer, Heidelberg, Dec. 2012. (Page 24.)
21. N. Bitansky and O. Paneth. ZAPs and non-interactive witness indistinguishability from indistinguishability obfuscation. In Y. Dodis and J. B. Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 401–427. Springer, Heidelberg, Mar. 2015. (Pages 3, 8, 9, and 10.)
22. D. Boneh and M. K. Franklin. Identity-based encryption from the Weil pairing. In J. Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 213–229. Springer, Heidelberg, Aug. 2001. (Page 1.)
23. D. Boneh, A. Sahai, and B. Waters. Functional encryption: Definitions and challenges. In Y. Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 253–273. Springer, Heidelberg, Mar. 2011. (Pages 1, 5, 6, 19, 21, and 22.)
24. E. Boyle, K.-M. Chung, and R. Pass. On extractability obfuscation. In Y. Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 52–73. Springer, Heidelberg, Feb. 2014. (Page 58.)
25. Z. Brakerski, I. Komargodski, and G. Segev. Multi-input functional encryption in the private-key setting: Stronger security from weaker assumptions. *Journal of Cryptology*, 31(2):434–520, Apr. 2018. (Page 58.)
26. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, Oct. 2001. (Pages 2, 57, and 58.)
27. P. Chaidos and G. Couteau. Efficient designated-verifier non-interactive zero-knowledge proofs of knowledge. In J. B. Nielsen and V. Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 193–221. Springer, Heidelberg, Apr. / May 2018. (Page 37.)
28. Y. Chen, V. Vaikuntanathan, B. Waters, H. Wee, and D. Wichs. Traitor-tracing from LWE made simple and attribute-based. In A. Beimel and S. Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 341–369. Springer, Heidelberg, Nov. 2018. (Page 37.)
29. J. Chotard, E. Dufour Sans, R. Gay, D. H. Phan, and D. Pointcheval. Decentralized multi-client functional encryption for inner product. In T. Peyrin and S. Galbraith, editors, *ASIACRYPT 2018, Part II*, volume 11273 of *LNCS*, pages 703–732. Springer, Heidelberg, Dec. 2018. (Page 58.)
30. J. Chotard, E. Dufour Sans, R. Gay, D. H. Phan, and D. Pointcheval. Multi-client functional encryption with repetition for inner product. Cryptology ePrint Archive, Report 2018/1021, 2018. <https://eprint.iacr.org/2018/1021>. (Page 58.)
31. G. Couteau and D. Hofheinz. Designated-verifier pseudorandom generators, and their applications. In Y. Ishai and V. Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 562–592. Springer, Heidelberg, May 2019. (Page 37.)
32. I. Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In J. Feigenbaum, editor, *CRYPTO’91*, volume 576 of *LNCS*, pages 445–456. Springer, Heidelberg, Aug. 1992. (Page 3.)
33. A. De Caro, V. Iovino, A. Jain, A. O’Neill, O. Paneth, and G. Persiano. On the achievability of simulation-based security for functional encryption. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 519–535. Springer, Heidelberg, Aug. 2013. (Page 1.)
34. D. Dolev, C. Dwork, and M. Naor. Nonmalleable cryptography. *SIAM Journal on Computing*, 30(2):391–437, 2000. (Page 3.)
35. A. Escala, G. Herold, E. Kiltz, C. Ràfols, and J. Villar. An algebraic framework for Diffie-Hellman assumptions. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 129–147. Springer, Heidelberg, Aug. 2013. (Page 7.)
36. P. Farshim, B. Libert, K. G. Paterson, and E. A. Quaglia. Robust encryption, revisited. In K. Kurosawa and G. Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 352–368. Springer, Heidelberg, Feb. / Mar. 2013. (Pages 3 and 4.)
37. L. Fortnow. The complexity of perfect zero-knowledge (extended abstract). In A. Aho, editor, *19th ACM STOC*, pages 204–209. ACM Press, May 1987. (Page 8.)

38. S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, Oct. 2013. (Page 58.)
39. S. Garg, C. Gentry, S. Halevi, and M. Zhandry. Functional encryption without obfuscation. In E. Kushilevitz and T. Malkin, editors, *TCC 2016-A, Part II*, volume 9563 of *LNCS*, pages 480–511. Springer, Heidelberg, Jan. 2016. (Page 58.)
40. R. Géraud, D. Naccache, and R. Rosie. Robust encryption, extended. In M. Matsui, editor, *CT-RSA 2019*, volume 11405 of *LNCS*, pages 149–168. Springer, Heidelberg, Mar. 2019. (Page 4.)
41. O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In A. Aho, editor, *19th ACM STOC*, pages 218–229. ACM Press, May 1987. (Page 8.)
42. S. Goldwasser, S. D. Gordon, V. Goyal, A. Jain, J. Katz, F.-H. Liu, A. Sahai, E. Shi, and H.-S. Zhou. Multi-input functional encryption. In P. Q. Nguyen and E. Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 578–602. Springer, Heidelberg, May 2014. (Page 58.)
43. S. Goldwasser, Y. T. Kalai, R. A. Popa, V. Vaikuntanathan, and N. Zeldovich. Reusable garbled circuits and succinct functional encryption. In D. Boneh, T. Roughgarden, and J. Feigenbaum, editors, *45th ACM STOC*, pages 555–564. ACM Press, June 2013. (Pages 37 and 58.)
44. S. Gorbunov, V. Vaikuntanathan, and H. Wee. Functional encryption with bounded collusions via multi-party computation. In R. Safavi-Naini and R. Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 162–179. Springer, Heidelberg, Aug. 2012. (Pages 1 and 58.)
45. V. Goyal, O. Pandey, A. Sahai, and B. Waters. Attribute-based encryption for fine-grained access control of encrypted data. In A. Juels, R. N. Wright, and S. De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 89–98. ACM Press, Oct. / Nov. 2006. Available as Cryptology ePrint Archive Report 2006/309. (Page 1.)
46. J. Groth. Simulation-sound NIZK proofs for a practical language and constant size group signatures. In X. Lai and K. Chen, editors, *ASIACRYPT 2006*, volume 4284 of *LNCS*, pages 444–459. Springer, Heidelberg, Dec. 2006. (Page 37.)
47. J. Groth, R. Ostrovsky, and A. Sahai. Non-interactive zaps and new techniques for NIZK. In C. Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 97–111. Springer, Heidelberg, Aug. 2006. (Pages 3, 8, 9, 10, and 37.)
48. V. Iovino and K. Zebrowski. Simulation-based secure functional encryption in the random oracle model. In K. E. Lauter and F. Rodríguez-Henríquez, editors, *LATINCRYPT 2015*, volume 9230 of *LNCS*, pages 21–39. Springer, Heidelberg, Aug. 2015. (Page 1.)
49. S. Katsumata, R. Nishimaki, S. Yamada, and T. Yamakawa. Designated verifier/prover and preprocessing NIZKs from Diffie-Hellman assumptions. In Y. Ishai and V. Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 622–651. Springer, Heidelberg, May 2019. (Page 37.)
50. J. Katz, A. Sahai, and B. Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In N. P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 146–162. Springer, Heidelberg, Apr. 2008. (Page 1.)
51. A. B. Lewko and B. Waters. Decentralizing attribute-based encryption. In K. G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 568–588. Springer, Heidelberg, May 2011. (Page 4.)
52. Y. Lindell. A simpler construction of cca2-secure public-key encryption under general assumptions. In E. Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 241–254. Springer, Heidelberg, May 2003. (Pages 36 and 44.)
53. C. Matt and U. Maurer. A definitional framework for functional encryption. In *IEEE 28th Computer Security Foundations Symposium, CSF 2015, Verona, Italy, 13-17 July, 2015*, pages 217–231, 2015. (Pages 1, 2, 6, 7, 11, 13, 21, 23, 41, 48, 49, 50, 62, and 63.)
54. M. Naor. On cryptographic assumptions and challenges (invited talk). In D. Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 96–109. Springer, Heidelberg, Aug. 2003. (Page 13.)
55. M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *22nd ACM STOC*, pages 427–437. ACM Press, May 1990. (Pages 3, 32, and 41.)
56. A. O’Neill. Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556, 2010. <http://eprint.iacr.org/2010/556>. (Page 1.)
57. C. Peikert and S. Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In A. Boldyreva and D. Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 89–114. Springer, Heidelberg, Aug. 2019. (Page 37.)
58. J. M. Pollard. Kangaroos, monopoly and discrete logarithms. *Journal of Cryptology*, 13(4):437–447, Sept. 2000. (Page 24.)
59. W. Quach, R. D. Rothblum, and D. Wichs. Reusable designated-verifier NIZKs for all NP from CDH. In Y. Ishai and V. Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 593–621. Springer, Heidelberg, May 2019. (Page 37.)

60. R. D. Rothblum, A. Sealfon, and K. Sotiraki. Towards non-interactive zero-knowledge for NP from LWE. In D. Lin and K. Sako, editors, *PKC 2019, Part II*, volume 11443 of *LNCS*, pages 472–503. Springer, Heidelberg, Apr. 2019. (Page 37.)
61. A. Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *40th FOCS*, pages 543–553. IEEE Computer Society Press, Oct. 1999. (Page 9.)
62. A. Sahai and B. R. Waters. Fuzzy identity-based encryption. In R. Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 457–473. Springer, Heidelberg, May 2005. (Page 1.)
63. N. Soroush, V. Iovino, A. Rial, P. B. Roenne, and P. Y. A. Ryan. Verifiable inner product encryption scheme. Cryptology ePrint Archive, Report 2020/122, 2020. <https://eprint.iacr.org/2020/122>. (Page 4.)

A Overview of the UC Framework

We use the universal composability (UC) framework introduced by Canetti [26] and provide a brief overview in this section. The goal of the UC framework is to capture what it means for a protocol to securely carry out a task. For this, we need to describe an ideal process and prove that no (efficient) environment can distinguish the real process and the ideal process, where the real-process is an execution of the protocol. Ideal processes are typically captured by ideal functionalities, which can be thought of as an incorruptible machine providing capabilities to different parties. These guarantees can depend on the corruption status of the parties in the system.

Protocol and protocol instances. Formally, a protocol π is an algorithm for a distributed system and formalized as an interactive Turing machine. An ITM has several tapes, for example an identity tape (read-only), an activation tape, or input/output tapes to pass values to its program and return values back to the caller. A machine also has a backdoor tape where (especially in the case of ideal functionalities) interaction with an adversary is possible or corruption messages are handled. While an ITM is a static object, UC defines the notion of an ITM instance (denoted ITI), which is defined by the extended identity $\text{eid} = (M, id)$, where M is the description of an ITM and $id = (\text{sid}, \text{pid})$ is a string consisting of a session identifier sid and a party identifier $\text{pid} \in \mathcal{P}$. An instance, also called a session, of a protocol π (represented as an ITM M_π) with respect to a session number sid is defined as a set of ITIs $\{(M_\pi, id_{\text{pid}})\}_{\text{pid} \in \mathcal{P}}$ where $id_{\text{pid}} = (\text{sid}, \text{pid})$.

The real process can now be defined by an environment \mathcal{Z} (a special ITI) that spawns exactly one session of the protocol in the presence of an adversary \mathcal{A} (also a special ITI), where \mathcal{A} is allowed to corrupt ITIs and gain their control. Which ITIs and in which form they can be corrupted is defined in a corruption model. In this work, we follow the static corruption model, which says that a party is either corrupted right from the beginning of the execution, or never. While static corruption is often needed when encryption schemes are involved, it also makes it possible to reason in a fine-grained fashion about the security of a system by looking at the specific set of corrupted parties. We note that this corruption set in the system is always known to the environment.

The output of the execution is the bit output by \mathcal{Z} and is denoted by $\text{EXEC}_{\pi, \mathcal{A}, \mathcal{Z}}(k, z, r)$ where k is the security parameter, $z \in \{0, 1\}^*$ is the input to the environment, and randomness r for the entire experiment. Let $\text{EXEC}_{\pi, \mathcal{A}, \mathcal{Z}}(k, z)$ denote the random variable obtained by choosing the randomness r uniformly at random and evaluating $\text{EXEC}_{\pi, \mathcal{A}, \mathcal{Z}}(k, z, r)$. Let $\text{EXEC}_{\pi, \mathcal{A}, \mathcal{Z}}$ denote the ensemble $\{\text{EXEC}_{\pi, \mathcal{A}, \mathcal{Z}}(k, z)\}_{k \in \mathbb{N}, z \in \{0, 1\}^*}$.

Ideal-world process The ideal process is formulated with respect to an ITM Func which is called an ideal functionality. In the ideal process, the environment \mathcal{Z} interacts with Func , an ideal-world adversary (often called the simulator) \mathcal{S} and a set of trivial, i.e., dummy ITMs representing the protocol machines that forward to the functionality whatever is provided as inputs to them by the environment (and return back whatever received from the functionality). In the ideal world, the ideal-world adversary (aka the simulator) can decide to corrupt parties. All corruptions are handled by the functionality which can assign more or less capabilities to the adversary depending on which parties are declared as corrupted in the system.

We denote the output of this ideal-world process by $\text{EXEC}_{\text{Func}, \mathcal{A}, \mathcal{Z}}(k, z, r)$ where the inputs are as in the real-world process. Let $\text{EXEC}_{\text{Func}, \mathcal{S}, \mathcal{Z}}(k, z)$ denote the random variable obtained by choosing the randomness r uniformly at random and evaluating $\text{EXEC}_{\text{Func}, \mathcal{S}, \mathcal{Z}}(k, z, r)$. Let $\text{EXEC}_{\text{Func}, \mathcal{S}, \mathcal{Z}}$ denote the ensemble $\{\text{EXEC}_{\text{Func}, \mathcal{S}, \mathcal{Z}}(k, z)\}_{k \in \mathbb{N}, z \in \{0, 1\}^*}$.

Hybrid worlds To model setup, the UC framework knows so-called hybrid worlds, which are worlds where the protocol under considerations invoke make use of ideal functionalities as subroutines (i.e., they invoke an ideal process as a subroutine). In this work, we use an authenticated repository and channel as assumed ideal functionalities.

Secure Realization and Composition In a nutshell, a protocol securely realizes an ideal functionality Func if the real-world process (where the protocol is executed) is indistinguishable from the ideal-world process (relative to Func):

Definition A.1. Let us denote by $\mathcal{X} = \{X(k, z)\}_{k \in \mathbb{N}, z \in \{0,1\}^*}$ and $\mathcal{Y} = \{Y(k, z)\}_{k \in \mathbb{N}, z \in \{0,1\}^*}$ two distribution ensembles over $\{0, 1\}$. We say that \mathcal{X} and \mathcal{Y} are indistinguishable if for any $c, d \in \mathbb{N}$ there exists a $k_0 \in \mathbb{N}$ such that $|\Pr[X(k, z) = 1] - \Pr[Y(k, z) = 1]| < k^{-c}$ for all $k > k_0$ and all $z \in \bigcup_{\kappa \leq k^d} \{0, 1\}^\kappa$. We use the shorthand notation $\mathcal{X} \approx \mathcal{Y}$ to denote two indistinguishable ensembles.

Definition A.2. Let Func be an ideal functionality and let π be a protocol. We say that π securely realizes Func if for any (efficient) adversary \mathcal{A} there exists an (efficient) ideal-world adversary (the simulator) \mathcal{S} such that for every (efficient) environment \mathcal{Z} it holds that $\text{EXEC}_{\pi, \mathcal{A}, \mathcal{Z}} \approx \text{EXEC}_{\text{Func}, \mathcal{S}, \mathcal{Z}}$, as defined above.

Note that the definition in [26] allows to capture in a more fine-grained way the context in which a protocol is executed as a further condition on the environment. We do not need this in our work and the statement holds for all contexts. The realization notion is composable, which is roughly speaking the guarantee that whenever in a certain context, the ideal process is used (e.g. as setup in a hybrid world) then it can be replaced by the protocol realizing it.

B Consistency for Different Types of Functional Encryption

In the present treatment we considered the standard public-key single-input FE setting. Nevertheless, our consistency notions can be relevant also when considering other FE settings. In particular it is easy to extend our treatment and results to give some meaningful guarantees for secret-key [9, 24, 38, 39, 43, 44] and multi-input/multi-client FE [10, 13, 25, 42]. It is worth noting that stronger (more tailored) notions of consistency might be conceivable in these cases, which depends on the intended applications. This may lead to a new modeling of these properties as an interesting further direction.

For secret-key FE, since the input provider and the setup generator is one party (e.g., consider medical record and a system that assigns different access rights to different doctors) strong input consistency seems to be the only reasonable formulation and is basically covered in the verifiable FE paper [12].

Regarding Multi-Client FE, our notion still ensures that it is not efficiently possible to output a ciphertext ct (now consisting of n components) such that ct would not be explainable by a vector of input values (x_1, \dots, x_n) given output values y_i (derived from ct) for functions f_i . However, a stronger notion could be derived and analyzed in the UC setting that captures consistency across the components of a ciphertext (while allowing also ciphertexts that can be “mixed”). In the Multi-client setting, our notion gives again similar guarantees as above where the master public key is a vector. In this setting, it must be impossible to generate a ciphertext that yields inconsistent output values in the sense that no input vector (x_1, \dots, x_n) would exist. As in the case of Multi-Client FE, it might be interesting to define stronger and more specific notions of consistency for this setting. Especially setup consistency remains important in the case of *Decentralized* Multi-Client FE [2, 3, 29, 30], where a part of Charlie’s task, i.e., the functional key generation, is distributed among the clients. Although these schemes are proven with respect to a passive adversary, as soon as moving to the active case, consistency, as we define it in this work, is needed.

C Details Of the UC Analysis

C.1 Assumed Functionalities

For completeness, we describe here the channels that we assume as setup in this work. The authenticated broadcast channel between sender S and receiver R leaks the message to the adversary. We assume that the functionality follows the standard UC corruption model, i.e., in case sender S is corrupted, the adversary can choose the message that is sent but cannot send different messages to different recipients. Note that since we consider static corruption, our channels are slightly simplified in that we do not have to capture the case that a message that has been input before would be changed before delivery (since either the sender is corrupted upon sending or he is never corrupted).

Functionality $\text{Func}_{\text{auth}}^{S, R_1, \dots, R_n}$

The functionality is parameterized by sender (extended) identity S and receiver (extended) identities R_i . The functionality initializes an empty array M .

- On input $(\text{SEND}, \text{sid}, m)$ from S , store $M \leftarrow M \parallel m$ and output (SENT, m) to every R_i .
- Upon input $(\text{GETMSGs}, \text{sid})$ from the adversary (on the backdoor tape) output M to the adversary.

Furthermore, we have point-to-point secure channels that are specified as follows:

Functionality $\text{Func}_{\text{sec}}^{S, R}$

The functionality is parameterized by sender S and receiver R . The functionality initializes an empty array M .

- On input $(\text{SEND}, \text{sid}, m)$ from S , store $M \leftarrow M \parallel \text{length}(m)$ and output (SENT, m) to R .
- Upon input $(\text{GETMSGs}, \text{sid})$ from the adversary (on the backdoor tape) output M to the adversary.

The assumed, real-world repository is defined as follows:

Functionality $\text{Func}_{\text{basic-rep}, \mathcal{C}}^{\mathbf{A}, \mathbf{B}, t}$

The functionality is parameterized by a set $\mathcal{C} \subseteq \{0, 1\}^*$ and the party identifiers $\mathbf{A}, \mathbf{B}_i, i \in [t]$, it interacts with. The functionality manages a lookup table M , which is initially empty.

Input. Upon receiving $(\text{WRITE}, \text{sid}, x)$ from a party with party-id \mathbf{A} do: If $c \in \mathcal{C}$, then compute handle $h \leftarrow \text{getHandle}$ and store $M[h] \leftarrow c$. Return h to the calling party.

Output. Upon receiving $(\text{READ}, \text{sid}, h)$ from a party with party-id $\mathbf{B}_i, i \in [t]$, do: If $M[h] = \perp$ then return **noData**. Otherwise, i.e., if $M[h] \in \mathcal{C}$, return $M[h]$ to the calling party.

C.2 Proof of the UC Realization (Theorem 7.2)

Proof. We start by proving the first part of the theorem and only then prove the necessary direction.

Consistency implies the UC realization. We first describe the simulator \mathcal{S} for the dummy UC adversary \mathcal{D} , which basically means that \mathcal{S} receives the instructions by the environment \mathcal{Z} . We are in the static corruption case and thus can structure the proof by a case distinction according to the actual corruption set in the system to obtain the detailed claims of the theorem for each case, which we cast as separate lemmata below.

Simulation with only a corrupted input provider. Upon receiving $(\text{SETUP}, \text{sid})$ from $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\mathbf{A}, \mathbf{B}, \mathcal{C}, t}$, the simulator \mathcal{S} executes $(\text{msk}, \text{mpk}) \leftarrow \text{Setup}()$ to obtain the master public and private key and provides mpk to the environment as the message received by the dishonest input provider and decryptors (or leaked by the authenticated broadcast channel). Upon receiving $(\text{ASSIGNED}, \text{sid}, f, i)$ from $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\mathbf{A}, \mathbf{B}, \mathcal{C}, t}$ then \mathcal{S} computes $F_0 \leftarrow F_0 \cup \{f\}$ (where F_0 is initially empty) and evaluates $\text{sk}_f \leftarrow \text{KeyGen}(\text{mpk}, \text{msk}, f)$ and provides sk_f to the environment when asked $(\text{READ}, \text{sid})$ (meant for the secure channel between the setup generator and some corrupted decryptor that receives the functional key for f). Finally, update $F \leftarrow F \cup \{(\text{sk}_f, f)\}$ (again F is initially empty).

When given the adversarial input $(\text{WRITE}, \text{sid}, \text{ct})$ (an input meant for the real-world repository), the simulator does the following: it computes $x \leftarrow \text{Extract}(\text{ct}, F)$ and outputs $(\text{WRITE}, \text{sid}, x)$ to $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\mathbf{A}, \mathbf{B}, \mathcal{C}, t}$ in the name of \mathbf{A} and returns the obtained handle to the environment.

Upon receiving $(\text{READ}, \text{sid}, h, f)$ from $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\text{A}, \text{B}, \text{C}, t}$, recall the ciphertext ct associated to h (previously input by the corrupted input provider). Then compute $y \leftarrow \text{Dec}(\text{mpk}, f, \text{sk}_f, \text{ct})$, run $x \leftarrow \text{Extract}(\text{ct}, F)$, and provide the input $(\text{READ}, \text{sid}, h, f, (x, y))$ to $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\text{A}, \text{B}, \text{C}, t}$.

If any receiver is corrupted, the adversary can instruct them to input $(\text{READ}, \text{sid}, h)$ to directly obtain the ciphertext associated to h . The simulated answer simply returns the previously input ciphertext ct for handle h .

Lemma C.1. *For any environment \mathcal{Z} (and dummy adversary) with non-negligible advantage in distinguishing the real and ideal worlds (w.r.t. simulator \mathcal{S} above) when only corrupting party A (and possibly a subset of the receivers), we give a reduction ρ_1 to construct an adversary $\mathcal{A} := \rho_1(\mathcal{Z})$ that violates (at least one of) input-consistency or Q-extractability with non-negligible probability.*

Proof. For the reduction to the input-consistency game, consider the following events defined in the real-world execution: Let E_1 denote the event that at any point in the execution of \mathcal{Z} (with the protocol and dummy adversary \mathcal{D}), there is a handle h such that the set of associated output values $y_i^{(h)}$ obtained by party B on queries (READ, h, f_i) are such that $\{x' \in \mathcal{M} \mid \forall i : f_i(x') = y_i^{(h)}\} = \{\}$. Further, let E_2 be the event that at any point in the execution, $\text{Extract}(\text{ct}, F) \notin \mathcal{X} \cup \{\perp\}$ but $\mathcal{Q}(F_0)$, where F and F_0 are the assigned functions and function-key pairs, respectively.

As long as $\neg E_1 \cap \neg E_2$ holds in the execution, the simulator \mathcal{S} executes the real-world view perfectly, since none of the instructions marked with (\star) and $(\star\star)$ must be ever executed in the ideal world and any value y returned to an honest decryptor is explainable by an element $x \in \mathcal{X} \cup \{\perp\}$ that fulfills $f(x) = y$ for the queried function $f \in \mathcal{F}$ (and thus also \diamond is never observed).

Hence, let $\mathcal{A} := \rho_1(\mathcal{Z})$ be the consistency adversary that internally runs \mathcal{Z} and emulates the real-world view towards \mathcal{Z} , i.e., upon any request by \mathcal{Z} , ρ_1 emulates the actions of the protocol when generating its replies to \mathcal{Z} . Note that such an emulation is possible with access to the honestly generated master public-key and with access to the key-generation oracle provided by the input-consistency game. Once event E_1 is observed, ρ_1 identifies the handle h that caused the event and outputs the associated ciphertext ct_h stored for handle h . If event E_2 is observed, then ρ_1 outputs the ciphertext ct provided by the environment \mathcal{Z} that provoked event E_2 . Note that \mathcal{A} is efficiently implementable by the assumption of (the efficiently implementable function) $\text{preMap}()$ which can be used to detect E_1 (and further events in the other cases).¹⁴

The proof for this case is concluded by observing that \mathcal{A} contradicts the theorem assumptions: we either have that Extract does not satisfy Q-extractability (event E_2), or that the evaluation of a set of functions f_1, \dots, f_n for a handle h returned values y_1, \dots, y_n that do not have a common explanatin (event E_1) which implies that $\text{in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract})$ returns 1 with the same probability as the event that \mathcal{Z} provokes event E_1 . \square

Simulation with a corrupted setup generator. If only the setup generator is corrupted, this means that all inputs provided to party A must define a valid base value $x \in \mathcal{X}$ and thus, upon read queries by an honest receiver/decryptor B_i only 2.(a) in the description of $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\text{A}, \text{B}, \text{C}, t}$ is triggered. Thus the only thing the simulator \mathcal{S} has to do in this simulation is to translate dishonest actions of party C when sending values towards the other two parties and controlling the influence on Alice actions. Thus, upon receiving $(\text{SEND}, \text{sid}, m)$ for the channel from C to A, \mathcal{S} just remembers that $\text{mpk}_1 \leftarrow m$ (and similarly for the message sent to the receivers which is defined to be mpk_2). For messages $(\text{SEND}, \text{sid}, m)$ for the channel from C to some honest receiver B_i , the simulator does the following: it first parses m as (sk, f) and if $f \notin \mathcal{F}$, it gives up activation. Otherwise, it encrypts a fixed message $\text{ct} \leftarrow \text{Enc}(\text{mpk}_1, \bar{m})$ and performs a trial decryption $y \leftarrow \text{Dec}(\text{mpk}_2, f, \text{sk}, \text{ct})$. If $y = \diamond$ then give up activation. Otherwise, output $(\text{ASSIGN}, \text{sid}, f, i)$ to $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\text{A}, \text{B}, \text{C}, t}$ to assign the function f to be available for the decryptor. Note that in case the decryptor is dishonest, simply simulate the receipt

¹⁴ We note that by picking one ciphertext at random, ρ_1 could avoid the dependence on $\text{preMap}()$ at the cost of obtaining a security loss. However, since in order to define the ideal UC functionality (which must be an efficient program) such an efficient map must exist, and since assuming it here yields more straightforward arguments, we rely on it throughout this proof.

of message m . Finally, whenever activated by $\text{Func}_{\text{Rep},(\mathcal{F}^+,f_0,\mathcal{Q})}^{\text{A,B,C},t}$ with a delayed output ($\text{WRITE}, \text{sid}, x$), the simulator performs a trial encryption $\text{ct} \leftarrow \text{Enc}(\text{mpk}_1, x)$ and sends ACK for this operation if and only if $\text{ct} \neq \text{err}$. Otherwise, the simulator activates the environment as next entity. Note that ct is remembered as the ciphertext for handle h (for element x) if some decryptors are dishonest (to properly simulate the real-world repository's reply to a dishonest decryptor's query ($\text{READ}, \text{sid}, h$)).

Lemma C.2. *For any environment \mathcal{Z} (and dummy adversary) with non-negligible advantage in distinguishing the real and ideal worlds (w.r.t. simulator \mathcal{S} above) when only corrupting party C (possibly alongside a subset of receivers), we give reductions ρ_2 and ρ_3 to construct adversaries $\mathcal{A}_i := \rho_i(\mathcal{Z})$ such that at least one of the \mathcal{A}_i violates setup-consistency with non-negligible probability.*

Proof. For this case, we first make a hybrid argument: consider the protocol π' , which is defined as $\pi_{\text{FE}}^{\text{A,B,C},t}$ but where party A provides its received master-public key mpk_1 to parties B_i via an additional covert broadcast channel and where parties B_i , already upon receiving a functional key (sk, f) , perform a trial decryption and rejects the key if $\text{Dec}(\text{mpk}_2, f, \text{sk}, \text{Enc}(\text{mpk}_1, \bar{m})) = \diamond$, where mpk_2 is the master public key sent from party C to parties B_i . We observe that protocol π' and $\pi_{\text{FE}}^{\text{A,B,C},t}$ have equivalent behaviors as long as the environment is not able to provide an input ($\text{WRITE}, \text{sid}, x$) to party A that provokes event E_3 defined by the condition that $\text{Dec}(\text{mpk}_2, f, \text{sk}, \text{Enc}(\text{mpk}_1, \bar{m})) = \diamond$ but $\text{Dec}(\text{mpk}_2, f, \text{sk}, \text{Enc}(\text{mpk}_1, x)) \neq \diamond$ (of course within an honest receiver/decryptor and where f and sk have been received together from party C). In case $\neg E_3$, the function f is never evaluated upon input (READ, h, f) (for any h) by protocol π' , whereas in $\pi_{\text{FE}}^{\text{A,B,C},t}$ it might. Hence, let $\mathcal{A} := \rho_2(\mathcal{Z})$ be the adversary for the setup consistency game defined as follows: ρ_2 internally runs \mathcal{Z} and emulates the execution of protocol $\pi_{\text{FE}}^{\text{A,B,C},t}$ towards \mathcal{Z} (by monolithically executing all required protocol steps) until event E_3 is observed. In this case, ρ_2 outputs $(\text{mpk}_1, \text{mpk}_2, \text{sk}, x, \bar{m})$, where x and sk are the values fulfilling the condition of event E_3 . ρ_2 wins $\text{set-CONS}^{\text{FE}}(1^\lambda, \mathcal{A})$ with the same probability as event E_3 in the execution with \mathcal{Z} .

For the final argument, we proceed with the same pattern. This time, let E_4 be the event that in an execution with π' , \mathcal{Z} provokes for some handle h that a query ($\text{READ}, \text{sid}, h, f$) to party B, following a write instruction ($\text{WRITE}, \text{sid}, x$) to party A that returned this handle h , yields an output value $y \neq f(x)$. Again, the simulation \mathcal{S} interacting with the repository $\text{Func}_{\text{Rep},(\mathcal{F}^+,f_0,\mathcal{Q})}^{\text{A,B,C},t}$ is a perfect simulation of π' : in both worlds, functions are assigned that pass the trial-decryption test, and all function evaluations, for some assigned function f , yield $f(x)$ as output as in this case only instruction $(\star\star\star)$ of the ideal-world repository is executed. Again, we can upper bound the distinguishing advantage of the real and ideal world by the probability that \mathcal{Z} provokes E_4 . The corresponding reduction $\mathcal{A} := \rho_3(\mathcal{Z})$ emulates a real-world execution towards \mathcal{Z} , where it mimics the protocol actions of the honest parties A and B_i . This includes the receipt of two message mpk_1 and mpk_2 for parties A and B_i , respectively. The reduction ρ_3 , once it detects event E_4 is provoked, can output $(\text{mpk}_1, \text{mpk}_2, \text{sk}, f, x, x)$, where (sk, f) is defined as the key function pair provoked event E_4 . Hence, \mathcal{A} achieves $\text{set-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}) = 1$ with at least the probability of \mathcal{Z} provoking E_4 . \square

Simulation with a corrupted input provider and setup generator. The simulator in this case needs to combine parts of the above two simulation strategies for maliciously generated setup parameters. That is, it defines mpk as the claimed master public key that party C sends to an honest party B_i . For the other messages ($\text{SEND}, \text{sid}, m$) for the channel from C to B_i , the simulator again parses it as a key-function pair (sk, f) and does the validity tests as above and in case $f \in \mathcal{F}$ and the trial decryption $\text{Dec}(\text{mpk}, f, \text{sk}, \text{Enc}(\text{mpk}, \bar{m}))$ (with respect to one master public key) does not yield \diamond , then output ($\text{ASSIGN}, \text{sid}, f, i$) to $\text{Func}_{\text{Rep},(\mathcal{F}^+,f_0,\mathcal{Q})}^{\text{A,B,C},t}$. Simulating a dishonest receiver is straightforward.

For adversarial inputs by party A, \mathcal{S} again computes $x \leftarrow \text{Extract}(\text{ct}, F)$ for the current set of key-function pairs, provides ($\text{WRITE}, \text{sid}, x$) to $\text{Func}_{\text{Rep},(\mathcal{F}^+,f_0,\mathcal{Q})}^{\text{A,B,C},t}$ in the name of A and returns the obtained handle to the environment.

Finally, upon receiving ($\text{READ}, \text{sid}, h, f$) from $\text{Func}_{\text{Rep},(\mathcal{F}^+,f_0,\mathcal{Q})}^{\text{A,B,C},t}$ (upon a reading instruction by honest party B_i), obtain the ciphertext ct associated to h , compute $y \leftarrow \text{Dec}(\text{mpk}, f, \text{sk}_f, \text{ct})$ and $x \leftarrow \text{Extract}(\text{ct}, F)$ and provide the input ($\text{READ}, \text{sid}, h, f, (x, y)$) to $\text{Func}_{\text{Rep},(\mathcal{F}^+,f_0,\mathcal{Q})}^{\text{A,B,C},t}$.

Lemma C.3. *For any environment \mathcal{Z} (and dummy adversary) with non-negligible advantage in distinguishing the real and ideal worlds (w.r.t. simulator \mathcal{S} above) when only corrupting parties \mathbf{A} and \mathbf{C} (and possibly alongside a subset of receivers), we give reductions ρ_4 and ρ_5 to construct adversaries $\mathcal{A}_i := \rho_i(\mathcal{Z})$ such that at least one of the \mathcal{A}_i violates input-consistency or Q-extractability with non-negligible probability.*

Proof. We again make a first hybrid step and consider the protocol π'' , where each party \mathbf{B}_i , upon receiving a functional key \mathbf{sk} together with its claimed function f , performs a trial decryption $\text{Dec}(\text{mpk}, f, \mathbf{sk}, \bar{\text{ct}})$, where $\bar{\text{ct}} \leftarrow \text{Enc}(\text{mpk}, \bar{m})$ (i.e., just with respect to the claimed master public key). Analogously to above, let E_5 be the event defined for an execution characterized by the condition that the environment provides a ciphertext ct and a secret key (\mathbf{sk}, f) to some honest party \mathbf{B}_i such that $\text{Dec}(\text{mpk}, f, \mathbf{sk}, \bar{\text{ct}}) \neq \diamond$ and $\text{Dec}(\text{mpk}, f, \mathbf{sk}, \text{ct}) = \diamond$. Again, $\pi_{\text{FE}}^{\mathbf{A}, \mathbf{B}, \mathbf{C}, t}$ and π'' have an identical behavior for any honest party \mathbf{B}_j until event E_5 is triggered. As above, this yields a reduction $\mathcal{A} := \rho_4(\mathcal{Z})$, which emulates party \mathbf{B} 's actions towards \mathcal{Z} and if \mathcal{Z} provokes E_5 , it outputs $(\text{mpk}, \bar{\text{ct}}, \text{ct}, \{\mathbf{sk}\})$, where the triple $\bar{\text{ct}}, \text{ct}$ and \mathbf{sk} are the values provided by \mathcal{Z} that trigger event E_5 .

The final reduction is obtained by defining, for an execution of \mathcal{Z} with π'' the two events E_6 and E_7 (analogous to E_1 and E_2 above): Let E_6 denote the event that at any point in the execution of \mathcal{Z} (with the protocol and dummy adversary \mathcal{D}), there is a handle h such that the set of associated output values $y_i^{(h)}$ obtained by some honest receiver/decryptor \mathbf{B}_i on queries (READ, h, f_i) are such that $\{x' \in \mathcal{M} \mid \forall i : f_i(x') = y_i^{(h)}\} = \{\}$. Further, let E_7 be the event that at any point in the execution, $\text{Extract}(\text{ct}, F) \notin \mathcal{X} \cup \{\perp\}$ but $Q(F_0)$, where F and F_0 are the accepted functions and function-key pairs by the union of honest receivers. As long as $\neg E_6 \cap \neg E_7$, the outputs generated by any honest party \mathbf{B}_i are the decrypted values computed by the simulator and thus computed as in the real-world execution and none of the instructions (\star) and $(\star\star)$ are executed in this case (this includes that \diamond is not computed by the simulator as a return value y in this case). The final adversary $\mathcal{A} := \rho_5(\mathcal{Z})$ for $\text{st-in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract})$ is now designed analogous to ρ_1 : Here, ρ_5 only emulates the honest receivers' actions towards an environment. When it detects event E_6 , it outputs $(\text{mpk}, \text{ct}, \text{ct}, F)$ where ct is the ciphertext that provoked the output and F the key-function pairs provided (and simulated) w.r.t. the union of honest receivers. When it detects E_7 , then ρ_5 outputs the ciphertext ct that was input by the environment and obtained handle h . Again, we see that the adversary \mathcal{A} contradicts the theorem assumptions, since it either wins $\text{st-in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}, \text{Extract})$ (event E_6) or it proves that Q-extractability does not hold (event E_7). \square

Simulation only with corrupted receivers. This case handles the scenario when the input provider and the setup generator are honest, and we have to argue anything that the union of dishonest receivers/decryptors can do in the real world—where they have access to all ciphertexts and received a set of secret keys—is simulatable in the ideal world, where we by definition only leak the information $f(x)$ if x is an input and f is an assigned function to one of the dishonest receivers in the corruption set.

Since our repository construction is functionally equivalent to the construction presented by Matt and Maurer [53], we inherit the security statement in this case by their statement: in particular, assume the algorithms $\mathcal{S}_1, \mathcal{S}_2$, and \mathcal{S}_3 guaranteed to exist by CFE security. The simulator acts as follows (where the union of corrupted receivers is simply seen as one “large corrupted decryptor” and treat it as the one corrupted party): it first simulates the public parameter by executing $(\text{mpk}, s) \leftarrow \mathcal{S}_1()$. Upon receiving $(\text{ASSIGNED}, \text{sid}, f, j)$ (for some index i) from $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\mathbf{A}, \mathbf{B}, \mathbf{C}, t}$ then \mathcal{S} computes $F_0 \leftarrow F_0 \cup \{f\}$ and outputs (READ, h_i, f) for each available label h_1, \dots, h_n in $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\mathbf{A}, \mathbf{B}, \mathbf{C}, t}$ to obtain the associated value y_i , and execute $\text{sk}_f \leftarrow \mathcal{S}_2(f, y_1, \dots, y_n)[[s]]$. This key is then output whenever the simulator must simulate the transmission of this key towards a dishonest receiver.

On input $(\text{READ}, \text{sid}, h)$ from a dishonest decryptor (expecting a real-world ciphertext), do the following: if a handle h has been generated (i.e., assigned to a value by $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\mathbf{A}, \mathbf{B}, \mathbf{C}, t}$) and a ciphertext ct_h has already been simulated, then return ct_h to the adversary. Otherwise, the ciphertext for this handle is simulated as follows: for all already assigned functions $f_i \in F_0$, ask (READ, h, f_i) to $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\mathbf{A}, \mathbf{B}, \mathbf{C}, t}$ to obtain all function values y_1, \dots, y_k for this handle (of the underlying input) and also y_0 , which is the output of the leakage function f_0 . Simulate (and internally store) the ciphertext $\text{ct}_h \leftarrow \mathcal{S}_3(y_0, y_1, \dots, y_k)[[s]]$ and return ct_h as the answer to the adversary.

The reduction to CFE security directly follows from [53, Lemma 4.2].

All parties honest. The remaining case is a straightforward simulation of the real-world view: if all parties are honest, then the simulator simply has to generate and output honest public parameters as above.

The necessary direction. To prove that the consistency requirements described by in-CONS, set-CONS, and st-in-CONS are also necessary, we show that if, for a given scheme FE and all algorithms Extract, there are adversaries $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ such that $\Pr[\text{in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}_1, \text{Extract})]$, $\Pr[\text{set-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}_2)]$, or $\Pr[\text{st-in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}_3, \text{Extract})]$ is non-negligible, then we can distinguish a real execution of $\pi_{\text{FE}}^{\text{A,B,C},t}$ from any ideal-world execution for $\text{Func}_{\text{Rep},(\mathcal{F}^+,f_0,\mathcal{Q})}^{\text{A,B,C},t}$ and an arbitrary simulator \mathcal{S}' . Since we all-quantify over Extract, for the sake of the argument, we choose it to be the constant function $\text{Extract} := \text{unknown}$. From the assumed adversaries for this choice of Extract, we construct environments \mathcal{Z}_1 and \mathcal{Z}_2 that distinguish the real and ideal worlds, the former for the case when only A is corrupted, and for the latter when only party C is corrupted. For the sake of the argument, we only have to consider the simple setting with one honest receiver/decryptor B_1 and one corrupted receiver B_2 .

Case in-CONS: \mathcal{Z}_1 internally runs \mathcal{A}_1 and answers its KeyGen-queries for functional keys corresponding to a given function f by providing the input (ASSIGN, sid, $f, 2$) and obtaining the corresponding functional key from obtaining the value of the secure channel from party C to B_2 since B_2 is corrupted. When \mathcal{A}_1 outputs a ciphertext ct, \mathcal{Z}_1 does the following: it corrupts party A and instructs it (all via the dummy adversary) to issue the write instruction (WRITE, sid, ct) (destined for the real-world repository) and expect handle h in return. Then issue n read instructions (READ, sid, h, f_i) to (honest) receiver B_1 to obtain n values $y_i, 1 \leq i \leq n$. If $\{x' \in \mathcal{M} \mid \forall i : f_i(x') = y_i\} = \{\}$ then \mathcal{Z}_1 outputs 1 and otherwise it outputs 0. It is clear that \mathcal{Z}_1 never outputs 1 when interacting with any simulator and functionality $\text{Func}_{\text{Rep},(\mathcal{F}^+,f_0,\mathcal{Q})}^{\text{A,B,C},t}$, since $\text{Func}_{\text{Rep},(\mathcal{F}^+,f_0,\mathcal{Q})}^{\text{A,B,C},t}$ will never output inconsistent values. On the other hand, \mathcal{Z}_1 outputs 1 whenever \mathcal{A}_1 detects an inconsistency. Hence, the distinguishing advantage equals $\Pr[\text{in-CONS}^{\text{FE}}(1^\lambda, \mathcal{A}_1, \text{Extract} := \text{unknown})]$.

Case set-CONS: \mathcal{Z}_2 internally runs \mathcal{A}_2 until it outputs $(\text{mpk}_1, \text{mpk}_2, \text{sk}, f, x_0, x_1)$. \mathcal{Z}_2 then interacts with honest party A and B_1 as follows via the party C that it corrupts: it instructs the corrupted party to send (via the channels) mpk_1 and mpk_2 to the respective parties A and B_1 and sk to B_1 . Then it chooses a bit j at random and provides the inputs (WRITE, sid, x_j) and (WRITE, sid, x_{1-j}) to party A to obtain the handles h_1 and h_2 . Finally, it provides the input (READ, h_1, f) to party B_1 . If it obtains an answer y_1 (and the input is hence not ignored) and $y_1 \neq f(x_j)$, then \mathcal{Z}_2 outputs 1 as its decision bit. Otherwise, it provides the input (READ, h_2, f) to party B_1 . If it obtains an answer y_2 (and the input is hence not ignored) and $y_2 \neq f(x_{1-j})$, then \mathcal{Z}_2 outputs 1 as its decision bit. Finally, if exactly one query returned an answer, then \mathcal{Z}_1 outputs 1 as its decision bit. In any other case, \mathcal{Z}_2 outputs 0. By assumption, the probability that at least one of the equation $\text{Dec}(\text{mpk}_2, f, \text{sk}, \text{Enc}(\text{mpk}_1, x_j)) \neq \diamond$ holds is at least $\Pr[\text{set-CONS}(1^\lambda, \mathcal{A}_2)]$ and therefore, with probability at least $\frac{\Pr[\text{set-CONS}(1^\lambda, \mathcal{A}_2)]}{2}$ reading h_1 will return a result when interacting with the protocol (as otherwise, the key sk would be ignored). When interacting with the ideal functionality and any simulator \mathcal{S}' , then by definition of $\text{Func}_{\text{Rep},(\mathcal{F}^+,f_0,\mathcal{Q})}^{\text{A,B,C},t}$, either both requests are ignored, or both requests return the expected result $f(x_j)$ and $f(x_{1-j})$ upon the first and second decryption, respectively. Therefore, we conclude that the probability that \mathcal{Z}_2 outputs 1 when interacting with the ideal world is zero and we obtain a distinguishing advantage of at least $\frac{\Pr[\text{set-CONS}(1^\lambda, \mathcal{A}_2)]}{2}$ for \mathcal{Z}_2 .

Case st-in-CONS: \mathcal{Z}_3 internally runs \mathcal{A}_3 which outputs $(\text{mpk}, \text{ct}_0, \text{ct}_1, \{(\text{sk}_1, f_1), \dots, (\text{sk}_n, f_n)\})$. \mathcal{Z}_3 then corrupts parties A and C and instructs party C to send mpk to (honest) party B_1 . \mathcal{Z}_3 then instructs party A to write each ciphertext ct_j to the real-world repository to obtain handle h_j . Then, it does the following for each secret key $\text{sk}_i, 1 \leq i \leq n$:

1. It instructs party C to send sk_i to party B_1 , followed by a query (READ, sid, h_1, f_i).
2. If the above requests got ignored, then it instructs party C to resend sk_i . In any case, it then issues (READ, sid, h_2, f_i).
3. When exactly one of the two queries gets ignored, then \mathcal{Z}_3 outputs decision bit 1 and halts. If both returned a value, it records them as $y_i^{\text{ct}_j}$ and proceeds with the next functional key.

If no decision has been reached, then \mathcal{Z}_3 defines the sets $S_j := \{x' \in \mathcal{M} \mid \forall i : f_i(x') = y_i^{(ct_j)}\}$ and outputs decision bit 1 if and only if at least one of S_0 or S_1 is equal to the empty set. Otherwise, \mathcal{Z}_3 outputs 0. When \mathcal{Z}_3 interacts with the protocol, then the values $y_i^{ct_j}$ are computed exactly as in the game $\text{st-in-CONS}(1^\lambda, \mathcal{A}_3, \text{Extract} := \text{unknown})$ and the decision bit is 1 if and only if the winning condition of the game is met (note that a query (w.r.t. f_i) is ignored if and only if the ciphertext decrypted to \diamond with respect to sk_i in the above execution with \mathcal{Z}_3). On the other hand, if \mathcal{Z}_3 is interacting with the ideal system, then it would never output 1, as $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\text{A,B,C,t}}$ does either answer both queries per evaluation for f_i or none. Furthermore, the above sets S_j are non-empty as ensured by $\text{Func}_{\text{Rep}, (\mathcal{F}^+, f_0, \mathcal{Q})}^{\text{A,B,C,t}}$ (ensured by instructions (\star) and $(\star\star)$). On the other hand, it outputs 1 when interacting with the real protocol if and only if the conditions of the game are fulfilled by the output of \mathcal{A}_3 , yielding a distinguishing advantage of $\Pr[\text{st-in-CONS}(1^\lambda, \mathcal{A}_3, \text{Extract} := \text{unknown})]$.

This concludes the proof of the theorem. □