

Analysis of Key Wrapping APIs: Generic Policies, Computational Security

Guillaume Scerri
University of Bristol
Bristol, United Kingdom
Email: guillaume.scerri@bristol.ac.uk

Ryan Stanley-Oakes
University of Bristol
Bristol, United Kingdom
Email: ryan.stanley@bristol.ac.uk

Abstract—We present an analysis of key wrapping APIs with generic policies. We prove that certain minimal conditions on policies are sufficient for keys to be indistinguishable from random in any execution of an API.

Our result captures a large class of API policies, including both the hierarchies on keys that are common in the scientific literature and the non-linear dependencies on keys used in PKCS#11. Indeed, we use our result to propose a secure refinement of PKCS#11, assuming that the attributes of keys are transmitted as authenticated associated data when wrapping and that there is an enforced separation between keys used for wrapping and keys used for other cryptographic purposes.

We use the Computationally Complete Symbolic Attacker developed by Bana and Comon. This model enables us to obtain computational guarantees using a simple proof with a high degree of modularity.

I. INTRODUCTION

A cryptographic API is an interface between a user and some trusted hardware, such as an HSM or a cryptographic token. These are deployed in insecure environments where the user’s machine (or the user herself) may be compromised. The cryptographic operations are performed inside the trusted hardware, keeping sensitive data such as secret keys out of reach of the host machine. One widely-used cryptographic API is PKCS#11, which is described in a cryptographic standards document consisting of hundreds of pages [OAS15a].

Many cryptographic APIs, including PKCS#11, allow key management commands such as key wrapping, the encryption of one key under another to facilitate secure key transport. Unfortunately, numerous key recovery attacks on PKCS#11 have been found, many of which exploit key wrapping commands that are explicitly authorised by the policy in the standard [Clu03], [DKS08], [DKS10]. Indeed, these attacks have been replicated on real cryptographic tokens that implement PKCS#11 [BCFS10].

Key management APIs are widely deployed, but the attacks on PKCS#11 demonstrate the dangers of not rigorously defining an API and its security properties. This led to the development of a number of formal models for key management APIs that have appeared in the scientific literature. However, these existing models have a number of important shortcomings that we describe below.

One weakness of the existing work is that proofs of API security typically depend on a particular choice of API security

policy. A key management API policy specifies how keys can be used. For example, the policy should determine which keys may be used to wrap a particular key. Many of the attacks on PKCS#11 exploit weaknesses in the policy. For example, the PKCS#11 policy allows keys to be used for both wrapping other keys and decrypting data. Therefore, an adversary can request the decryption of a wrap and receive a sensitive key in the clear. Nevertheless, there has so far been no *generic* analysis of key management policies themselves, where one gives conditions on the policy, rather than describing it in full, and proves that these conditions are sufficient to prevent attacks. Instead, previous security proofs for API designs have specified a particular policy for their particular design and proved some security property of the API with this policy [CC09], [CS09], [KSW11]. There have been type-based analyses of APIs that allow for more general statements, but these works only prove security against key recovery (not that keys are indistinguishable from random) and do not consider the effect of key corruption [CFL13], [AFL13].

Not analysing generic policies is a serious shortcoming of existing work: poor policy design results in devastating attacks that are easy to implement, but there are no general results on what constitutes a valid policy. Proprietary APIs are likely to implement their own policies, not the particular ones that have been used in API designs in the academic literature. Providing a generic framework for evaluating policies gives the ability to analyse both the API designs already in use and any future designs that might appear.

A different kind of policy, not considered at all in previous analyses of key management APIs, is a high-level key management policy, independent of the choice of key management API, that must be enforced by any API. For example, consider the use of cryptographic tokens in a large enterprise, where employees have a position in a hierarchy. Suppose the enterprise sees employees near the bottom of the hierarchy as more likely to be compromised than those higher up, so the high-level key management policy on keys is simply that the compromise of keys for employees lower down the hierarchy should not compromise keys higher up. To enforce this, the API policy would need to prevent keys of lower-ranked employees being used to wrap keys of higher-ranked employees. We call the high-level key management policy the *enterprise security policy* to reflect this example,

but in general an enterprise policy could be a complicated, non-linear relationship between different components of the organisation¹. Since previous analyses in the literature make no link between API policies and enterprise policies, it is unclear how a security property proved for a key management API is sufficient for the security needs of the users of the API.

Another weakness of existing work is that security proofs are not modular. Analyses of key management APIs typically make specific assumptions about the cryptographic mechanisms used by the API, and the security proofs only hold under these assumptions. For example, while Kremer, Steel and Warinschi give an API design with strong computational security guarantees, their proof relies on the assumption that key wrapping is deterministic and that wraps are unforgeable and indistinguishable from random strings [KSW11]. It is not clear how making different assumptions on the wrapping mechanism would affect the proof. The same can be said for the API design by Cachin and Chandran, whose security definition, and hence the proof of security, is very closely tied to the design of the API and the specific cryptographic security notions assumed for the primitives used by the API [CC09]. Kremer, Künnemann and Steel also make restrictive assumptions about the wrapping primitive used by their key management functionality, namely that it is deterministic and has Key-Dependent Message (KDM) security [KKS13]. All of these results only verify the security of a particular key wrapping API with a particular policy and a particular wrapping mechanism. Additionally, the proof in the paper by Kremer, Künnemann and Steel is in the GNUC framework. This means that the security notion used for key management should compose naturally with the security of other functionalities offered by a cryptographic API, but this composability comes at the cost of having a complex model and a proof that is hard to verify.

The Cortier-Steel and Cortier-Steel-Wiedling APIs allow arbitrary terms to be processed by the API and only have one type of encryption for, in particular, both keys and data [CS09], [CSW12]. As input, their APIs can take a term representing the encryption of a key, together with some data encrypted under another key. This is intended for running complex cryptographic protocols without ever having to give the value of keys to an untrusted machine. This is a desirable design for cryptographic APIs, however it is far from what is implemented in practice (in particular, what is implemented in PKCS#11). Our approach is closer to what is implemented in practice, where terms are typically either the wrapping of one key or the encryption of some data.

The Cortier-Steel API uses a very rich policy that describes relationships between agents, but this policy is expressible using our simple definition of policy. The Cortier-Steel-Wiedling

¹ As another example, suppose an organisation has a CEO (A), a CFO (B), an HR Manager (C) and a Product Designer (D). D is lower-ranked than B and C, who are equally-ranked and lower-ranked than A. If C's machine is compromised, and hence C's secret keys are compromised, then one might expect the company's employee data to be compromised, but not the company's financial data. Both (respectively, neither) sets of data are expected to be compromised if A (resp., D) is compromised.

API additionally attaches timestamps to keys so that their validity can change over time. This allows for an API that recovers after corruption. We do not consider the issue of time here, and leave such an extension as future work. Finally, both these works perform their proofs in the Dolev-Yao model. In this respect the computational guarantees we give in this paper are strictly stronger.

We improve on existing work in a number of ways: we avoid specifying a particular API policy; we use an approach that is modular with respect to cryptographic assumptions; we prove strong, computational security guarantees with a proof that is easy to verify and we parameterise our definition of API security by the security needs of the users of the API. Ultimately, our result is significantly more general, and our proof is significantly more modular, than in previous analyses of key management APIs. We can verify the security, with respect to different enterprise security policies, of a wide variety of key wrapping APIs with different policies and different wrapping mechanisms.

In the next subsection we outline our results.

A. Our Contribution

We view a cryptographic API as the composition of a *key wrapping API* that carries out key wrapping, and an *external API* that carries out other cryptographic operations. All keys stored by the API are called using a public handle, and each handle has an associated *attribute* that is either *internal* or *external*, referring to the intended use of the key pointed to by the handle. The key wrapping API will reject calls to wrap under keys with external attributes, and the external API will reject calls to carry out the other cryptographic operations using keys with internal attributes.

Our security goal for a key wrapping API is that external keys are indistinguishable from freshly generated keys, even after being wrapped and unwrapped using internal keys. As we exemplify in Section VII, this notion of security for key wrapping APIs is sufficient to preserve the security of the other cryptographic operations carried out by the external API.

We remark that indistinguishability from random is impossible to achieve for any key that can be used for wrapping² so the key wrapping API will need to enforce the separation between internal and external keys. While we assume the API does not explicitly allow attributes to be changed from external to internal, it is possible that an adversary could implicitly change the attribute of a key by wrapping and unwrapping it, giving it a new handle and possibly a new attribute. It will therefore be necessary to use a secure AE-AD scheme to wrap keys, so that the wrapping mechanism binds the attributes of the wrapped key to the ciphertext.

From now on, unless stated otherwise, we will use the term API to refer to the key wrapping API inside a larger cryptographic API, since this is the main focus of our analysis.

²There is a trivial distinguishing attack in this case: the distinguisher creates a wrap under the true key, then attempts to decrypt this wrap using its challenge key.

An API has a *security policy*, or simply *policy*, that determines whether one key can be used to wrap another. This policy is not simply how the API ought to behave, but rather the actual rules for wrapping implemented in the code of the API.

We analyse a *generic* key wrapping API in which the security policy is not fixed. That is, rather than choosing a particular policy as is done elsewhere in the literature, we leave it underspecified and find conditions on the policy that are sufficient for secure key management. While previous works conflate two roles of policies - defining the actual internal *behaviour* of the API and defining what properties are *expected* from the API - we clearly separate these roles into the API security policy and the enterprise policy, respectively, and determine when the former is sufficient to satisfy the latter.

Furthermore, our API design does not insist on a particular wrapping mechanism. Unlike other API designs, the wrapping mechanism can be implemented in a variety of ways without affecting our security theorem or its proof - in particular, both deterministic and randomised wrapping mechanisms are supported. As with the policy, we specify conditions on the security of the mechanism, capturing confidentiality and integrity, that are sufficient for the security of the API.

Our approach gives strong, computational guarantees but with a simple proof in the symbolic model that holds under a number of different computational assumptions.

Now we detail some of the key technical aspects of our result.

STRONG SECURITY GUARANTEES. We prove that certain minimal conditions on the policy and the wrapping mechanism are sufficient to guarantee the secrecy of external keys. We prove security in a strong, cryptographic sense, namely that external keys are indistinguishable from random keys, even in the presence of powerful Probabilistic Polynomial-time Turing machine (PPT) adversaries. We additionally prove that composing a secure key wrapping API with a secure encryption scheme results in a secure cryptographic API where users are able to wrap keys and encrypt data. As an application, we propose a secure refinement of PKCS#11, forbidding certain attribute combinations and forcing the wrap mechanism to be a secure AE-AD scheme that correctly transmits the attributes of wrapped keys.³

GENERIC POLICIES. We assume the existence of an enterprise security policy and an API security policy but do not fix either one. The enterprise policy is specified as relationships between the attributes of keys. Intuitively, these relationships determine how the compromise of keys will propagate. Then we give conditions that say whether or not the security policy of an API is valid with respect to the enterprise security policy; simply separating internal and external keys is not enough for security. For example, suppose the enterprise security policy says that the key k_1 is of a higher security level than k_2 . Then a valid API security policy will not allow k_1 to be wrapped

under k_2 since, if k_2 is compromised, a wrap of k_1 under k_2 is insecure, leading to the trivial compromise of k_1 .

MODULAR PROOF. We use the Computationally-Complete Symbolic Attacker for equivalence properties (CCSA), developed by Bana and Comon in 2014 [BC14]. This model uses symbolic formalism to prove meaningful, computational security statements. Proofs in this model are easy to verify, especially compared to the cryptographic proofs like those found in [CC09], [KSW11] and [KKS13]. In CCSA one expresses assumptions (on the policy and the wrapping mechanism) as axioms in first-order logic and deduces the security property from these axioms. This means that the theorem holds under any computational assumptions such that the axioms hold, giving a high degree of modularity.

Informally, our cryptographic axioms say that genuine wraps of keys are indistinguishable from wraps of fresh, random keys and that wraps are unforgeable. We prove that our axioms are sound if the wrapping scheme is built on an Authenticated Encryption with Associated Data (AE-AD) encryption scheme; this is a standard requirement in modern, symmetric cryptography. However, we also prove that the axioms are sound when wrapping satisfies a *deterministic* variant of this notion (i.e. not using random nonces). Since our proof depends only on the axioms, not on the computational assumptions, we can obtain results about APIs that use completely different wrapping mechanisms without any changes to the proof. One of our conditions for an API policy to be valid is that it forbids the creation of key cycles. However, if we assume KDM security for the wrapping scheme as in [KKS13], then we can relax this condition on the policy (see Remark 5). This is in contrast to existing works, where relaxing the assumption on key cycles or the assumption that encryption is randomised would require substantial changes to the security proof.

AUGMENTING THE CCSA FRAMEWORK. In previous works, CCSA has only been applied to cryptographic protocols. Ours is the first use of CCSA to analyse a system that is outside the scope of the original paper on CCSA. To prove our substantial result, we required novel, computationally-sound axioms. One such new axiom captures how a proof can be split into disjoint cases. Our new axioms can be used in security proofs outside the key management setting without needing to reprove their computational soundness. In this respect we have augmented the CCSA model.

Another recent work has developed CCSA axioms for various cryptographic primitives in order to tackle a larger class of cryptographic protocols [BC16]. Their work independently developed axioms for manipulating branching and ‘if’ statements in the CCSA model; most of the core axioms related to these constructions are similar in our work and in theirs. The main difference between our core axioms and those developed in [BC16] is that our case disjunction axiom is strictly stronger than theirs, at the cost of a more involved soundness proof. Our stronger case disjunction axiom is required to deal with the very rich branching that arises in the context of key-management.

³We are aware that this is likely to be far from what is currently implemented on real-world cryptographic tokens running PKCS#11-compliant APIs. Our result should be viewed simply as guidance for future token designs.

CAVEATS. We prove security when the number of API queries made by the adversary is arbitrary but, because of our reliance on CCSA, this number must be independent of the security parameter used by the underlying cryptographic primitives. Nevertheless, all of the attacks found on real APIs either use a fixed number of queries, as captured by our model, or are attacks on weak cryptographic primitives (such as in [BFK⁺12]) for which our axioms would not be sound.

We accept that our separation of external and internal keys is difficult to enforce in practice; in particular, the attributes of keys must be securely bound to the keys when wrapping (such as with AE-AD) so that external and internal keys can never be confused. However, this separation enforces a standard industry practice, as recommended by NIST: “a single key should be used for only one purpose (e.g., encryption, authentication, key wrapping, random number generation, or digital signatures)” [NIS12]. Indeed, most of the API designs in the literature enforce this separation [CC09], [KSW11], [CS09], [CSW12].

We remark that our composition theorem in Section VII only applies to cryptographic APIs with wrap, unwrap, encrypt, decrypt and corrupt actions, rather than arbitrary cryptographic primitives. However, the proof is largely independent of the specific cryptographic game for encryption and would therefore be easy to adapt to other primitives.

II. KEY WRAPPING APIS

In this Section we define the execution model for a key wrapping API and the security one should expect of such an API.

A. Execution Model

We assume the existence of the sets \mathcal{K} of keys, \mathcal{H} of handles and \mathcal{D} of attributes (data). The set \mathcal{D} has a particular subset \mathcal{E} of external attributes. A wrapping mechanism wm consists of the triple $(\text{keygen}, \text{wrap}, \text{unwrap})$ of algorithms. The algorithm keygen takes a security parameter η (in unary) as input and returns an element of \mathcal{K} . We assume that $\text{keygen}(1^\eta)$ always returns a key of length $\text{keylen}(1^\eta)$. The algorithm wrap takes as input a key $k \in \mathcal{K}$, an attribute $a \in \mathcal{D}$ and a second key $k' \in \mathcal{K}$, and returns $\text{wrap}(k, a, k')$, the wrap packet of the key k with attribute a under the key k' . We do not specify how the wrap packet depends on the attribute of the wrapped key, nor how it uses any randomness. The algorithm unwrap takes a wrap packet and a key as input and returns a key and an attribute. Both wrap and unwrap can have access to the security parameter, if necessary. We assume that all wrapping mechanisms are correct, that is, for all keys $k, k' \in \mathcal{K}$ and all attributes $a \in \mathcal{D}$, $\text{unwrap}(\text{wrap}(k, a, k'), k') = (k, a)$.

A security policy \mathbb{P} is an algorithm that takes two attributes as input and returns a bit. We make this choice since a policy ought not to depend on the values of the keys themselves (otherwise it could leak unintended information about the keys and would have to be evaluated by the secure hardware, not the API). We remark that a more general wrapping policy could

use the value of a global clock as an additional argument (to handle key lifecycles and key revocation, etc.), but we leave this for future work.

Definition 1. A key wrapping API API is a program parameterised by the tuple $(\mathcal{K}, \mathcal{D}, \mathcal{E}, \mathcal{H}, \text{wm}, \mathbb{P})$. The API maintains, in its state st , a map $\text{st.val} : \mathcal{H} \rightarrow (\mathcal{K} \cup \{\zeta\}) \times (\mathcal{D} \cup \{\zeta\})$ that records the association between handles, which are public names, and pairs consisting of a key and its attribute. In what follows, we specify how an adversary interacts with an API.

An adversary, interacting with the API, can request wraps and unwraps via the handles of keys. If $\text{st.val}(h) = (\zeta, \zeta)$ then we say h is unused in state st . There is also a function freshhdl that takes a state st as input and returns a handle $\text{freshhdl}(\text{st})$ that is unused in st .

Let st_0 be an initial state of the API. Then the initial configuration of the API is the map val encoded by st_0 . An initial configuration is called *honest* if, for all handles h , $\text{st}_0.\text{val}(h)$ is either not initialised or initialised with an output of $\text{keygen}(1^\eta)$. Intuitively, this corresponds to a device with honestly generated keys already present. From now on, all key wrapping APIs are assumed to have honest initial configurations. This assumption is reflected in the generation of the initial state of APIs in the experiment used to define security.

Note that we assume that each key has a single attribute. We will also assume the API does not explicitly allow the user to change the attribute of a key. These assumptions are without loss of generality, since any attributes that can be easily changed by the adversary cannot be used to preserve meaningful security and so are omitted from this discussion. Note that any “attributes” irrelevant to security can be encoded in the state and we simply do not allow the policy to depend on these. Furthermore, if a key could be generated without an attribute, then the attribute would need to be fixed before the key could be used and so we merely assume for convenience that this decision takes place before the key is generated.

Intuitively, if a key’s attribute belongs to \mathcal{E} then that key is intended to protect data, as opposed to other keys. Our security goal will be that external keys (keys with attributes in \mathcal{E}) remain indistinguishable from random after being managed by the API and therefore are ideal for cryptographic use. Note that this goal is impossible to achieve for any keys used for key wrapping, since an adversary can distinguish the real key from a random key by attempting to unwrap with its test key.

Even though separating external and wrapping keys might seem restrictive, as soon as a wrapping key can also be used for decryption, keys can be recovered by the adversary (as demonstrated in [BCFS10]). Therefore such a separation must be enforced by any secure policy. While we only consider here the wrapping module used by cryptographic APIs, in Section VII we give a formal argument that a secure key wrapping API can be composed with an encryption scheme without undermining the security of the encryption scheme, assuming this separation of key roles.

There are three actions that a user can perform in its interaction with a key wrapping API: wrapping, unwrapping and corruption. Obviously, the corruption action is not intended to be implemented on real APIs, but is used here to reason about the security of APIs in the presence of an adversary who can obtain the values of particular keys, for example through side-channel attacks.

When a key is corrupted, we add its attribute to a list of corrupted attributes. We settle for recording corrupted attributes rather than corrupted keys as, by construction, every key with the same attribute has the same capabilities. In other words, if one corrupts a key with attribute a , every key that can be wrapped by a key with attribute a will be compromised, irrespective of exactly which key was corrupted. It is therefore enough to log that *some* key with attribute a has been corrupted. As a consequence, if one wants to distinguish the corruption of two keys, these keys should be given different attributes.

A *fully-specified action* is an action (wrap, unwrap or corrupt), together with the handles of the keys relevant to the action. Fully-specified actions define the execution of the API. Formally, if \mathbf{A} is a fully-specified action of the API, then \mathbf{A} is an element of one of the sets $\{\mathbf{W}(h_1, h_2) \mid h_1, h_2 \in \mathcal{H}\}$, $\{\mathbf{U}(h) \mid h \in \mathcal{H}\}$ or $\{\mathbf{C}(h) \mid h \in \mathcal{H}\}$. Obviously the unwrap action will require an additional argument corresponding to the ciphertext to be unwrapped, but the value of this argument does not determine the overall structure of the API execution (as it could depend on some random coins) and so it is not considered part of the fully-specified action $\mathbf{U}(h)$.

- If $\mathbf{A} = \mathbf{W}(h_1, h_2)$, then the API computes $(k_1, a_1) \leftarrow \text{st.val}(h_1)$ and $(k_2, a_2) \leftarrow \text{st.val}(h_2)$. If $k_1, k_2 \neq \perp$ and $\mathbb{P}(a_1, a_2) = 1$, then the API returns $\text{wrap}(k_1, a_1, k_2)$. Otherwise, the API returns \perp .
- If $\mathbf{A} = \mathbf{U}(h_1)$, then the API takes an additional input x from the user, computes $(k_1, a_1) \leftarrow \text{st.val}(h_1)$ and $(k_2, a_2) \leftarrow \text{unwrap}(x, k_1)$. If $k_2, a_2 \neq \perp$ and $\mathbb{P}(a_2, a_1) = 1$, then $\text{st.val}(\text{freshhdl}(\text{st})) \leftarrow (k_2, a_2)$. Otherwise, $\text{st.val}(\text{freshhdl}(\text{st})) \leftarrow (\perp, \text{bad})$ where bad is a particular attribute (used to denote handles that are not fresh, but do not point to a key). In either case, $\text{freshhdl}(\text{st})$ is returned to the user.
- If $\mathbf{A} = \mathbf{C}(h_1)$, then the user is asking to learn the value of the key pointed to by h_1 . The API computes $(k_1, a_1) \leftarrow \text{st.val}(h_1)$, does not carry out any checks, and returns k_1 . By definition, if h is unused in st then the API returns \perp . For our security property, the state maintains a list of *corrupt attributes*. Therefore, for this action, the API updates its state as follows: $\text{st.cor} \leftarrow \{a_1\} \cup \text{st.cor}$.

B. Security

We consider three objectives of an adversary in its execution of a key wrapping API. First, an adversary may try to learn (part of) the value of an external key, in order to compromise a cryptographic scheme using this key in the wider cryptographic API. Second, the adversary may try to change the attributes of a key (external or otherwise), in order

to circumvent the security policy. Third, the adversary may try to import its own keys in order to wrap with these keys. The second and third objectives can be combined: the adversary succeeds in either objective if there is a handle pointing to a key, attribute pair that was not pointed to by a handle in the initial configuration of the API (either because the key is new, or the attribute of an honestly generated key has changed).

We consider an API to be *secure* if and only if no adversary can achieve these objectives with *uncompromised* handles. That is, if a handle is uncompromised and it points to an external key, then that key cannot be distinguished from a random key; we call this the *key secrecy* property. Furthermore, any uncompromised handle must point to a key, attribute pair from the initial configuration of the API; we call this the *handle consistency* property.

What is meant by calling a handle compromised (or uncompromised) is that, for example, an organisation might consider keys assigned to lower-status employees compromised if the keys of senior management are compromised, but not vice-versa. This is an example of an *enterprise security policy*. So, given the list of attributes of keys lost via the corrupt action, the enterprise security policy determines which handles are compromised according to their attributes. We formalise this below.

In general, we model an enterprise security policy by a *sacrifice function* $\text{sacr} : \mathcal{D} \rightarrow \mathcal{P}(\mathcal{D})$. Intuitively, the set $\text{sacr}(a)$ should be thought of as the set of attributes of keys compromised as a result of compromising a key with attribute a . This has nothing to do with the use of the API; it is a pre-existing relationship between attributes that the API needs to respect, by not allowing the compromise of keys to propagate beyond what is specified by the sacrifice function.

We assume without loss of generality that $a \in \text{sacr}(a)$ for all $a \in \mathcal{D}$ and that sacr is *transitive* in the following sense: if $a \in \text{sacr}(b)$ and $b \in \text{sacr}(c)$ then $a \in \text{sacr}(c)$.

The sacrifice function is derived from the existing relationships between users of the API. The example enterprise security policy given above, where compromise propagates down the employee hierarchy (but not up), would be captured by giving an ordering $<$ on attributes and putting $\text{sacr}(a) = \{a' \in \mathcal{D} \mid a' \leq a\}$ for all $a \in \mathcal{D}$. A more refined enterprise security policy would separate the organisation into departments and say that the compromise of keys must not propagate across different departments. In this case, one would partition the set of attributes into classes (\mathcal{D}_i) and, for each i , put $\text{sacr}(a) = \{a' \in \mathcal{D}_i \mid a' \leq a\}$ for all $a \in \mathcal{D}_i$.

There could be a very conservative enterprise security policy, where the compromise of a key does not compromise any other keys. In this case, each key would have a unique attribute and $\text{sacr}(a) = \{a\}$ for each attribute a . We will see that the only APIs secure for this sacrifice function are ones that disable key wrapping altogether.

Now we formalise the notion of *compromise*, which follows easily from the sacrifice function and the list of corrupt attributes.

Definition 2. If $a \in \text{st.cor}$ then we say a is a *corrupt attribute* in state st . If a is corrupt in state st and $a' \in \text{sacr}(a)$, then we say a' is a *compromised attribute* in state st .

If $\text{st.val}(h) = (k, a)$ and a is a corrupt (respectively, compromised) attribute in state st , then we say h is a *corrupt handle* (respectively, *compromised handle*) in state st .

We have two security experiments, capturing the secrecy of uncompromised keys and the consistency of uncompromised attributes.

Our security experiments take an adversary \mathcal{A} , integers m and n and attributes a_1, \dots, a_n as input. The handles h_1, \dots, h_n are initialised with randomly generated keys and the attributes a_1, \dots, a_n . The integer m is the exact number of oracle queries that \mathcal{A} makes in either experiment.

In both security experiments, the adversary has access to three oracles capturing the wrap, unwrap and corrupt actions of the API. There is also a fourth *test* oracle, which may only be queried once, whose behaviour depends on a bit b . The adversary submits a handle and, if $b = 0$, the oracle returns the key pointed to by this handle. If $b = 1$, a new key is randomly generated and returned to the adversary. When performing this test query, the API checks that the attribute of the submitted handle is *external* and has not been compromised. After the test query has been made, the API checks that any corrupt query does not compromise the tested handle.

If the tested handle is compromised at the end of the key secrecy experiment, or its attribute is not external, then the output of this experiment is a random bit. Otherwise, the output of the experiment is the output of the adversary.

In the handle consistency experiment, the output of the experiment is 1 if, once the adversary has given its output, the state of the API contains an uncompromised handle that points to a key, attribute pair not in the initial configuration of the API. Otherwise, the output of the experiment is 0.

The formal description of the security experiments are given in Figures 1 and 2.

<p>Oracle $\mathcal{O}^{\text{wrap}}(x_1, x_2)$:</p> <p>$\text{st}.C \leftarrow \text{st}.C + 1$ If $\text{st}.C > m$ Return ζ $(k, a) \leftarrow \text{st.val}(x_1)$ $(k', a') \leftarrow \text{st.val}(x_2)$ If $k, k' \neq \zeta$ and $\mathbb{P}(a, a')$ Return $\text{wrap}(k, a, k')$ Return ζ</p>	<p>Oracle $\mathcal{O}^{\text{corrupt}}(x_1)$:</p> <p>$\text{st}.C \leftarrow \text{st}.C + 1$ If $\text{st}.C > m$ Return ζ $(k, a) \leftarrow \text{st.val}(x_1)$ If $\text{st}.T \in \text{sacr}(a)$ $\text{st}.V \leftarrow \perp$ $\text{st.cor} \leftarrow \{a\} \cup \text{st.cor}$ Return k</p>
<p>Oracle $\mathcal{O}^{\text{unwrap}}(x_1, x_2)$:</p> <p>$\text{st}.C \leftarrow \text{st}.C + 1$ If $\text{st}.C > m$ Return ζ $(k, a) \leftarrow \text{st.val}(x_1)$ $(k', a') \leftarrow \text{unwrap}(x_2, k)$ $h \leftarrow \text{freshhdl}(\text{st})$ $\text{st}.H \leftarrow \text{st}.H \cup \{h\}$ If $k' \neq \zeta$ and $\mathbb{P}(a', a)$ $\text{st.val}(h) \leftarrow (k', a')$ Else $\text{st.val}(h) \leftarrow (\zeta, \text{bad})$ Return h</p>	<p>Oracle $\mathcal{O}_b^{\text{test}}(x_1)$:</p> <p>If $\text{st}.C > m$ or $\text{st}.T \neq \perp$ Return ζ $(k, a) \leftarrow \text{st.val}(x_1)$ $k' \leftarrow \text{keygen}(1^\eta)$ If $a \notin \mathcal{E}$ then $\text{st}.V \leftarrow \perp$ For $c \in \text{st.cor}$ If $a \in \text{sacr}(c)$ then $\text{st}.V \leftarrow \perp$ $\text{st}.T \leftarrow a$ If $b = 0$ Return k Return k'</p>

Figure 1. Wrap, Unwrap, Corrupt and Test Oracles

Experiment $\text{Exp}_b^{\text{KEYSEC}}(\mathcal{A}, a_1, \dots, a_n, m)$:

$\text{st}.C \leftarrow 0$
 $\text{st}.T \leftarrow \perp$
 $\text{st}.V \leftarrow \top$
 $\text{st.cor} \leftarrow \emptyset$
For $1 \leq i \leq n$
 $k_i \leftarrow \text{keygen}(1^\eta)$
 $\text{st.val}(h_i) \leftarrow (k_i, a_i)$
 $\text{st}.H \leftarrow \{h_1, \dots, h_n\}$
 $\mathcal{O}_b \leftarrow (\mathcal{O}^{\text{wrap}}, \mathcal{O}^{\text{unwrap}}, \mathcal{O}^{\text{corrupt}}, \mathcal{O}_b^{\text{test}})$
 $b' \leftarrow \mathcal{A}^{\mathcal{O}_b}(1^\eta)$
If $\text{st}.V = \top$ Return b'
 $b'' \leftarrow \text{keygen}(1^\eta)$
Return b''

Experiment $\text{Exp}^{\text{HDLCON}}(\mathcal{A}, a_1, \dots, a_n, m)$:

$b \leftarrow 1$
 $\text{st}.C \leftarrow 0$
 $\text{st}.T \leftarrow \perp$
 $\text{st}.V \leftarrow \top$
 $\text{st.cor} \leftarrow \emptyset$
 $\text{st.val}_0 \leftarrow \emptyset$
For $1 \leq i \leq n$
 $k_i \leftarrow \text{keygen}(1^\eta)$
 $\text{st.val}(h_i) \leftarrow (k_i, a_i)$
 $\text{st.val}_0 \leftarrow \text{st.val}_0 \cup \{(k_i, a_i)\}$
 $\text{st}.H \leftarrow \{h_1, \dots, h_n\}$
 $\mathcal{O}_b \leftarrow (\mathcal{O}^{\text{wrap}}, \mathcal{O}^{\text{unwrap}}, \mathcal{O}^{\text{corrupt}}, \mathcal{O}_b^{\text{test}})$
 $x \leftarrow \mathcal{A}^{\mathcal{O}_b}(1^\eta)$
For $h \in \text{st}.H$
 $(k, a) \leftarrow \text{st.val}(h)$
If $(k, a) \notin \text{st.val}_0$
If for all $a_c \in \text{st.cor}$, $a \notin \text{sacr}(a_c)$
Return 1
Return 0

Figure 2. Key Secrecy Experiment b and Handle Consistency Experiment

Definition 3. Let $\text{API} = (\mathcal{K}, \mathcal{D}, \mathcal{E}, \mathcal{H}, \text{wm}, \mathbb{P})$. We say the API is secure if, for all integers m and n , all $\vec{a} \in \mathcal{D}^n$ and all polynomial-time adversaries \mathcal{A} , the following advantages are both negligible functions of η :

$$\text{Adv}_{\text{API}^m}^{\text{KEYSEC}}(\mathcal{A}) := \mathbb{P}_b [\text{Exp}_b^{\text{KEYSEC}}(\mathcal{A}, \vec{a}, m) = b] - \frac{1}{2}$$

$$\text{Adv}_{\text{API}^m}^{\text{HDLCON}}(\mathcal{A}) := \mathbb{P} [\text{Exp}^{\text{HDLCON}}(\mathcal{A}, \vec{a}, m) = 1]$$

III. STATEMENT OF THE MAIN THEOREM

In this Section we define the class of *valid* API security policies, recall the definition of a secure AE-AD scheme and state the main theorem of the paper: if a key wrapping API uses a secure AE-AD scheme for its wrapping mechanism and its security policy is valid, then the API is secure.

There are three conditions required of the API policy. First, we require that if a key with attribute a can be wrapped under a key with attribute a' , then the enterprise security policy should say that the compromise of a key with attribute a' compromises keys with attribute a . This is a reasonable assumption as any wrap under a key with attribute a' will no longer provide any security for the key with attribute a .

Second, we require that the policy forbids wrapping under external keys. This enforces the separation between keys used by the API and keys protected by the API.

Finally, we require that valid policies forbid the creation of key cycles. This is necessary, since standard security notions for encryption do not imply security in the presence of key cycles (see e.g. [CGH12]).

Definition 4. Let \mathbb{P} be an API security policy. We say \mathbb{P} is *valid* with respect to the sacrifice function sacr if:

- For all $a, a' \in \mathcal{D}$, if $\mathbb{P}(a, a') = 1$, then $a \in \text{sacr}(a')$.
- For all $a \in \mathcal{D}$ and $e \in \mathcal{E}$, $\mathbb{P}(a, e) = 0$.
- The *policy graph*, i.e. the graph on \mathcal{D} where there is an edge $a \rightarrow a'$ if and only if $\mathbb{P}(a', a) = 1$ (if and only if a key with attribute a can be used to wrap a key with attribute a'), is acyclic.

Remark 5. This definition of a valid policy is minimal since, if any of the conditions are not met, then attacks will be possible:

- 1) If the first condition is false, then it will be possible to create a wrap of a key k such that corrupting the wrapping key does not compromise the handle pointing to k . So k will be an uncompromised key (according to the enterprise policy) that is trivially distinguishable from random.
- 2) If the second condition is false, then there is an external key k such that the adversary can create a wrap under k and then try to decrypt the wrap with its challenge key.
- 3) The third condition is necessary since standard notions of encryption security do not imply security in the presence of key cycles. We could remove this condition and show that the key wrapping APIs are still secure, if we assume KDM security for the wrapping mechanism. In this case, we would simply weaken the constraint on the secrecy axiom to permit key cycles (see Section V), prove the soundness of this axiom (by slightly modifying the soundness proof of the original axiom) and then the proof of the symbolic security property from the axioms (Section VI) would need only minor changes to accommodate the new constraint.

Even though the above definition captures a large class of sensible security policies, it is all we need in order to prove that a key wrapping API is a secure API, provided that the wrapping mechanism is built using either a (randomised) *secure AE-AD scheme*, as defined in [Rog02] or a Deterministic Authenticated Encryption scheme as in [RS06]. We refer to these as randomised and deterministic AE-AD schemes, respectively. We recall these security notions here:

A triple $\Pi = (\text{keygen}, \text{enc}, \text{dec})$ is called a randomised (resp. deterministic) *Secure Authenticated Encryption with Associated Data (AE-AD) scheme* if:

- Π is *correct*, i.e. for all messages m , data a , nonces r and all $k \leftarrow_s \text{keygen}(1^n)$, $\text{dec}(\text{enc}(m, a, r, k), a, r, k) = m$ (where nonces are ignored in the deterministic case),
- Π is *private*, i.e. no PPT adversary can distinguish between an encryption oracle and an oracle returning random strings of the appropriate length,
- Π is *authenticated*, i.e. given access to an encryption oracle, no PPT adversary may produce triples (c, a, r) (or pairs (c, a) in the deterministic case) that decrypt successfully but where the c was not output by the encryption oracle.

Full definitions are provided in Appendix ??.

Now we state the main Theorem of this paper.

Theorem 1. Let $\text{API} = (\mathcal{K}, \mathcal{D}, \mathcal{E}, \mathcal{H}, \text{wm}, \mathbb{P})$ be a key wrapping API and let $\Pi = (\text{keygen}, \text{enc}, \text{dec})$ be a secure deterministic or randomised AE-AD scheme. Suppose $\text{wm} = (\text{keygen}, \text{wrap}, \text{unwrap})$ where wrap , on input (k, a, k') generates a fresh nonce r and returns the wrap packet $(\text{enc}(k, a, r, k'), a, r)$ and unwrap , on input (w, k) , parses w as (c, a, r) and returns $\text{dec}(c, a, r, k)$. Then, if \mathbb{P} is valid with respect to sacr , the key wrapping API is secure.

Remark 6. The hypotheses of this theorem are minimal. By Remark 5, a valid policy is necessary to prevent attacks. Furthermore, without assuming AE-AD security for the wrapping mechanism, or some other way of securely binding attributes to keys when wrapping, the first two attacks given in Remark 5 are still possible: the adversary simply wraps a key and unwraps it with new attributes, circumventing any restrictions given by the policy.

In order to prove Theorem 1, we define *symbolic* APIs, using the language of CCSA, in Section IV. We give axioms in Section V and prove these axioms sound under the assumptions of Theorem 1. In particular we give axioms that are sound for both secure deterministic and randomised AE-AD schemes. Therefore, due to the computational soundness of CCSA, if one can show that a symbolic security property is implied by the axioms, then the corresponding computational security property holds in any computational API where the wrapping mechanism is as in Theorem 1.

In Section VI, we show that the axioms entail the symbolic security properties corresponding to the security definition given in Definition 3. This proves Theorem 1.

IV. SYMBOLIC MODEL AND SOUNDNESS

To prove Theorem 1, we cast our problem in the setting of the computationally complete symbolic attacker model (CCSA) from [BC14]. The CCSA allows one to abstract away a lot of the complexity of the usual computational models, yet it entails computational guarantees.

In CCSA, we model the capabilities of an adversary as a list of first-order axioms that cannot be broken by any PPT adversary with non-negligible probability. For example, a secrecy property for encryption could be expressed as the formula $\{x\}_k \sim \{y\}_k$, meaning that no PPT adversary can distinguish an encryption of x from an encryption of y . This is in contrast to Dolev-Yao style symbolic models where one has to fully specify the abilities of the attacker as a finite list of inference rules.

A security property in CCSA is expressed as the indistinguishability of two lists of terms, which is itself written as a formula in first-order logic. For example, if ϕ is the list of terms output to the adversary by the API where h points to (k, a) in the initial state, and ϕ' is where h points to (k', a) , then the key secrecy property could be written $\phi \sim \phi'$. One then proves that the axioms entail the security property. In other words, an adversary can perform any action that does

⁴In fact, our key secrecy property is more complicated than this, since we need to specify that h is uncompromised.

not contradict the axioms, but still cannot break the security property.

Assigning a Turing Machine to each of the function symbols used in a symbolic API is called a *computational interpretation* of the symbolic API. For example, the function symbol $\{_ \}$ could be interpreted as the encryption algorithm enc . A *computational model* of the symbolic API is a computational interpretation where certain function symbols have fixed interpretations. For example, in a computational model, the function symbol EQ is interpreted as the machine that outputs 1 if its two inputs are equal (and outputs 0 otherwise). Crucially, the predicate symbol \sim is interpreted as computational indistinguishability of bitstrings. That is, the formula $x_1, \dots, x_n \sim y_1, \dots, y_n$ is interpreted as the statement that no polynomial-time adversary can distinguish the interpretations of x_1, \dots, x_n from the interpretations of y_1, \dots, y_n .

In a computational model, each axiom will be interpreted as a statement about Turing Machines. If the statement is true, then we say the axiom is *sound* in this model. Let A be a set of axioms. The computational soundness theorem for CCSA says that if $A \rightarrow \phi \sim \phi'$ then, in any computational model where the axioms in A are sound, the computational interpretations of the terms ϕ and ϕ' are indistinguishable: no polynomial-time adversary can distinguish them. In an API, the sequences of terms ϕ and ϕ' output by the API will depend on the inputs of the (active) distinguishing adversary, and these inputs are modelled using free function symbols g_i ; one must prove that $\phi \sim \phi'$ regardless of the interpretation of the symbols g_i .

The axiomatic approach of CCSA gives an inherent modularity to the model. In our case, we are able to specify only the properties of the API security policy that are necessary, rather than describing the whole policy, since these properties correspond to the axioms used in the proof. Thus, at no extra cost, we prove the security of multiple instantiations of key wrapping APIs.

A. Symbolic Execution

In CCSA, protocols⁵ have a fixed structure: it is known *a priori* what messages are expected at what stage (and therefore what checks to carry out before returning a message, and so on). In our computational API, even though the number of oracle queries is constant in the security parameter, the adversary may query its oracles in any order. Each possible order of queries will induce a different sequence of output terms from the symbolic API. The order does matter, since the inputs from the adversary (in particular, the inputs to the unwrap oracle) depend on the previous outputs seen by the adversary so far. Therefore we have to fix a sequence of oracle queries, obtaining a *restricted computational API*, and reason about the output terms of the symbolic API corresponding to this restricted computational API. Fixing a list of actions is

⁵The authors of [BC14] use the term protocols, but we will use their model to describe APIs. Having fixed a list of API oracle queries, we obtain a transition system that behaves like a protocol as described in [BC14]. Therefore we use ‘protocol’ and ‘API’ interchangeably.

not a significant restriction; we show in Appendix B how the computational security definition given in Section II reduces to the security of every possible restricted computational API.

We now give the language needed to describe the symbolic API.

Terms are built over the sets \mathcal{F} of function symbols, \mathcal{G} of free adversarial function symbols and \mathcal{N} of names. The set of all terms is written $\mathcal{T} := \mathcal{T}(\mathcal{F}, \mathcal{G}, \mathcal{N})$. Syntactic equality of the terms t and t' is written $t \equiv t'$. Every term has a sort. Within the sort message, there are subsorts key, nonce, longnonce, handle, attribute and bool. Names have sort key, nonce, longnonce, handle or attribute. There is a particular name bad of sort attribute. The set \mathcal{F} is described in Figure 3.

<ul style="list-style-type: none"> • Constants: true, false, fail • Conditional branching: if $_$ then $_$ else $_$ • Logical operators: $_ \wedge _$, $_ \vee _$, $\neg _$ • Pairing and projections: $(_, _)$, $\pi_1(_)$, $\pi_2(_)$, $\pi_3(_)$ • Encryption and decryption: $\text{enc}(_, _, _)$, $\text{dec}(_, _, _)$ • Equality test: $\text{wf}(_)$, $\text{EQ}(_, _)$ • Random string: $\\$(_, _)$ • Policy check: $\text{P}(_, _)$
--

Figure 3. Functions

We remark that certain function symbols above, e.g. the boolean operators and the projections have obvious intended meanings. We formalise these intended meanings either as fixed computational interpretations in Section IV-C, or give assumptions on these interpretations, which are reflected as axioms, in Section V. This means that for functions like encryption that have more than one obvious implementation, we permit any implementation satisfying the assumptions.

We introduce some abbreviations for ease of notation. We will write $[b]x : y$ for (if b then x else y) and $[b]x$ for (if b then x else fail). This notation is extended to lists of terms in the natural way: if $\vec{v} \equiv v_1, \dots, v_n$ and $\vec{w} \equiv w_1, \dots, w_n$, then $[b]\vec{v} : \vec{w} \equiv [b]v_1 : w_1, \dots, [b]v_n : w_n$.

We write $\{m\}_{k,a}^r \equiv (\text{enc}(m, a, r, k), a, r)$. This is just for convenience: when unwrapping, the adversary must submit a ciphertext, some associated data and a nonce to the API and we prefer to express this triple as a single term called the *wrap packet*. In the case of deterministic encryption, the nonce will simply be ignored. Furthermore, we write

$$\begin{aligned} \text{triple}(x) &\equiv \text{EQ}(x, (\pi_1(x), \pi_2(x), \pi_3(x))) \\ \text{uwp}(x, k) &\equiv [\text{triple}(x)] \text{dec}(\pi_1(x), \pi_2(x), \pi_3(x), k) \end{aligned}$$

This is so, when the adversary submits a wrap packet to uwp , one checks that this packet is a triple (x', x'', x''') and, if so, its components are taken for the inputs to dec . We use $\text{wf}(x) \equiv \neg \text{EQ}(x, \text{fail})$ to test whether terms (in particular, ciphertexts) are *well-formed*.

The set \mathcal{G} consists of infinitely many function symbols $g_i : \text{message}^i \rightarrow \text{message}$ for each $i \geq 0$, corresponding to the computation of the adversary on the terms it has seen previously.

Subterms are defined naturally as the internal components of terms, and $\text{st}(x)$ denotes the set of subterms of x . *Positions* of subterms within terms are defined as usual.

Formulae may be built over terms using predicate symbols. There is a predicate symbol \sim^i for each natural number i ; each may be used between lists of i terms. We abbreviate $x_1, \dots, x_i \sim^i y_1, \dots, y_i$ by $x_1, \dots, x_i \sim y_1, \dots, y_i$. We use the symbols $\neg, \vee, \wedge, \rightarrow$ to construct compound formulae as normal, but these should not be confused with the *function symbols* used to operate on terms of sort `bool`. For example, $b \rightarrow b'$ is a term of sort `bool`, whereas $(b \sim b') \wedge (b' \sim b'') \rightarrow (b \sim b'')$ is a (compound) formula.

We assume that each state q of a symbolic API is labelled by the fully-specified actions used to reach it from the initial state q_0 . Then, given an *initial configuration* $\text{val}_{q_0} : \mathcal{H} \rightarrow (\mathcal{K} \cup \{\text{fail}\}) \times (\mathcal{D} \cup \{\text{fail}\})$, we can recover the function $\text{val}_q : \mathcal{H} \rightarrow \text{message}$ by modifying val_{q_0} in the manner described below. For convenience of notation, we write $\mathbf{k}_q(h)$ and $\mathbf{a}_q(h)$ for $\pi_1(\text{val}_q(h))$ and $\pi_2(\text{val}_q(h))$ respectively.

We remark that val_q is an inherent property of the state q of the API and *not* a function symbol to be interpreted in a model: $\text{val}_q(h)$ will always be instantiated by a particular term of sort `message`, intended to be a key, attribute pair.

Finally we recall that, upon successfully unwrapping, the adversary is given a handle that points (via `val`) to the result of the unwrap. The new handle is determined by a function `freshhdl` that takes a state as input and returns a handle that is *unused* in state q , in the sense that $\text{val}_q(\text{freshhdl}(q)) \equiv (\text{fail}, \text{fail})$. Again, this function depends on the internal state of the API (in particular on val_q) and not on the interpretation: in each state q , `freshhdl`(q) will be instantiated by a particular handle and the `freshhdl` symbol will not appear in the terms output to the adversary.

Security properties in CCSA are expressed as the indistinguishability of sets of terms. We use the language introduced above to describe the sets of terms that will correspond to our key secrecy and handle consistency properties.

First, assume we are given a list of fully-specified actions $\mathbf{A}_1, \dots, \mathbf{A}_m$. This list induces a state transition system $q_0 \rightarrow q_1 \rightarrow \dots \rightarrow q_m$ as follows:

- The adversary's initial knowledge, ϕ_0 , is the set of pairs $\{(h_j, \mathbf{a}_{q_0}(h_j)) \mid 1 \leq j \leq m\}$.
- In state q_i , the adversary has executed the actions $\mathbf{A}_1, \dots, \mathbf{A}_{i-1}$ and received the list of terms ϕ_i . Then $\phi_{i+1} := \phi_i; s$ where s is the output from the next action \mathbf{A}_i , described below.

If $\mathbf{A}_i = \mathbf{W}(h_1, h_2)$, then

$$s \equiv [\mathbf{P}(\mathbf{a}_{q_i}(h_1), \mathbf{a}_{q_i}(h_2))] \{ \mathbf{k}_{q_i}(h_1) \}_{\mathbf{k}_{q_i}(h_2), \mathbf{a}_{q_i}(h_1)}$$

where r , a name of sort `nonce` used to indicate the (fresh) randomness of the wrapping scheme, occurs in no other output terms. The state q_{i+1} is such that $\text{val}_{q_{i+1}} = \text{val}_{q_i}$.

If $\mathbf{A}_i = \mathbf{U}(h)$, then

$$s \equiv [\text{wf}(u) \wedge \mathbf{P}(\pi_2(g_i(\phi_i)), \mathbf{a}_{q_i}(h))] \text{freshhdl}(q_i)$$

where $u \equiv \text{uwp}(g_i(\phi_i), \mathbf{k}_{q_i}(h))$. The state is updated as follows: if $h' \neq \text{freshhdl}(q_i)$ then $\text{val}_{q_{i+1}}(h') \equiv \text{val}_{q_i}(h')$

and, with u as above,

$$\begin{aligned} \text{val}_{q_{i+1}}(\text{freshhdl}(q_i)) & \equiv [\text{wf}(u) \wedge \mathbf{P}(\pi_2(g_i(\phi_i)), \mathbf{a}_{q_i}(h))] \\ & (u, \pi_2(g_i(\phi_i))) : (\text{fail}, \text{bad}), \end{aligned}$$

where $\text{bad} \in \mathcal{D}$ is a fixed attribute such that $\text{sacr}(\text{bad}) = \{\text{bad}\}$.

If $\mathbf{A}_i = \mathbf{C}(h)$, then $s \equiv \mathbf{k}_{q_i}(h)$ and q_{i+1} is such that $\text{val}_{q_{i+1}} = \text{val}_{q_i}$. Note that if h does not point to a key in state q_i , then $s \equiv \text{fail}$.

We remark that these output terms correspond to the API execution model in Section II.

B. Symbolic Security

In order to define the symbolic security properties, we define the symbolic analogues to the notions of *corrupt* and *compromised* attributes and handles introduced in Section II. We remark that, if the function symbols `EQ`, `∨` and `∧` are interpreted in the obvious way (which we will later insist on), the following symbolic definitions are equivalent to the computational definitions given previously.

Definition 7. Let $q \in \mathbf{Q}$. Let $\mathbf{C}(\mathcal{H})$ be the set of handles h such that the action $\mathbf{C}(h)$ occurred in the transitions $q_0 \rightarrow \dots \rightarrow q$. Then let $x \in \mathcal{T}$. We say x is a *corrupt attribute* in state q if $\text{cor}_q(x) \sim \text{true}$, where

$$\text{cor}_q(x) \equiv \text{EQ}(x, \text{bad}) \vee \bigvee_{h \in \mathbf{C}(\mathcal{H})} \text{EQ}(x, \mathbf{a}_q(h)).$$

Then for any $y \in \mathcal{T} \setminus \mathcal{H}$ we say y is a *compromised attribute* in state q if $\text{comp}_q(y) \sim \text{true}$, where

$$\text{comp}_q(y) \equiv \bigvee_{a \in \mathcal{D}} \bigvee_{b \in \text{sacr}(a)} \text{cor}_q(a) \wedge \text{EQ}(y, b).$$

Recall that $a \in \text{sacr}(a)$ for any $a \in \mathcal{D}$ and, if $b \in \text{sacr}(a)$ and $c \in \text{sacr}(b)$, then $c \in \text{sacr}(a)$. With the axiom $\text{EQ}(x, x) \sim \text{true}$ (which is obviously computationally sound for any term x if `EQ` is interpreted as equality of bitstrings), we can show that if a is a corrupt attribute in q , then a is a compromised attribute in q and if a is compromised in q and $b \in \text{sacr}(a)$, then b is compromised in q .

Definition 8. Let $q \in \mathbf{Q}$ and $h \in \mathcal{H}$. We say h is a *compromised handle* in state q and write $\text{comp}_q(h)$ if and only if $\mathbf{a}_q(h)$ is a compromised attribute in q .

By convention, `bad` is a corrupt attribute in any state. One might wonder why the attribute of `freshhdl`(q) becomes the *always compromised* attribute `bad` instead of remaining as `fail` in the event that the unwrap action fails. This is so that, if one tries to unwrap using a compromised handle, the new handle is considered compromised regardless of the properties of the adversary's arbitrary input $g_i(\phi_i)$ (which will typically depend on a random tape). This is mostly a technicality needed for uniformity later on: we will use the axioms to show that if a handle is *uncompromised* then it points to either a genuine

key, attribute pair or the pair (fail, fail), whereas if the handle used to unwrap is compromised then the axioms cannot be used to express $\text{val}_{q'}(\text{freshhdl}(q))$ in such a simple way. We can now state the symbolic version of security, and hence the symbolic version of our main Theorem. This Theorem states that that, given a handle h^* pointing to an *external* key k^* , the output ϕ^* of an API execution in which h^* is not compromised is indistinguishable from the output of the same execution but in which h^* points to a key k^\dagger not appearing in ϕ^* . Moreover this is true even if the adversary is given k^* in plaintext at the start of both executions. Additionally, we require that any uncompromised handle points to an honestly generated key together with the attribute that key had in the initial state.

Theorem 2. *Let $\text{API} = (\mathcal{K}, \mathcal{D}, \mathcal{E}, \mathcal{H}, \text{wm}, \text{P})$ be a key management API such that P is valid with respect to **sacr** and the axioms of Section V are sound. Say $\text{val}_{q_0^*}(h^*) = (k^*, a^*)$ where $a^* \in \mathcal{E}$. Let q_0^\dagger be identical to q_0^* except that $\text{val}_{q_0^\dagger}(h^*) = (k^\dagger, a^*)$, where k^\dagger is a key not appearing in ϕ^* (in particular $k^\dagger \neq k^*$). Finally, let t be a sequence of fully-specified actions inducing the API executions $(q_0^*, (k^*, \phi_0)) \rightarrow \dots \rightarrow (q^*, \phi^*)$ and $(q_0^\dagger, (k^*, \phi_0)) \rightarrow \dots \rightarrow (q^\dagger, \phi^\dagger)$. Then the following security properties hold:*

- 1) *key secrecy:* $[-\text{comp}_{q^*}(h^*)] \phi^* \sim [-\text{comp}_{q^\dagger}(h^*)] \phi^\dagger$
- 2) *handle consistency:*

$$\phi^\dagger, -\text{comp}_{q^\dagger}(h) \rightarrow \bigvee_{h_0 \in \mathcal{H}} \text{EQ}(\text{val}_{q^\dagger}(h), \text{val}_{q_0^\dagger}(h_0)) \sim \phi^\dagger, \text{true}$$

In order to relate this symbolic theorem to the computational security properties, we recall the computational interpretation of the logic from [BC14].

C. Computational Interpretation

To express the computational semantics of functions and predicates, we fix the computational interpretation of some functions. For example, the function EQ is supposed to express semantic equality, but without any restrictions it could be interpreted by any Turing machine and it would be effectively impossible to reason about it. Therefore we give fixed interpretations for certain symbolic objects and define a *computational model* to be any model satisfying our restrictions. Crucially, there is a great deal of freedom in how a computational model interprets the adversarial function symbols $g_i \in \mathcal{G}$, so that all realistic attacks (i.e. those that run in polynomial-time and do not have access to the internal randomness of the API) can be captured by computational models. Once these interpretation of terms are defined, the interpretation of the formula $\vec{x} \sim \vec{y}$ is simply that no PPT adversary can distinguish between the interpretation of \vec{x} and the interpretation of \vec{y} .

Intuitively, we ensure that the encryption and decryption symbols are interpreted as (possibly deterministic) authenticated encryption and decryption algorithms and we also enforce some straightforward restrictions on pairs, projections,

equality and boolean symbols. Details can be found in Appendix A. If the computational interpretation of a formula holds for any interpretation (with the aforementioned restrictions) of the function symbols, we say that this formula is *sound*.

Given these restrictions on the computational interpretation, we show that Theorem 2 entails Theorem 1. First we show that we can reduce the security game to the case of a single list of actions and then further reduce to when the test handle is chosen at the start of the execution. Details are provided in Appendix B.

V. AXIOMS

As in [BC14], we define axiom schemes $\gamma \parallel \theta(\vec{v})$ where γ is a *constraint* and \vec{v} is the vector of free variables in the formula θ . This is the set of all instances $\theta(\vec{v})\sigma$ where σ is an assignment of the free variables \vec{v} to ground terms, and is such that $\sigma \models \gamma$. For example, we define a constraint $\text{fresh}(r; \vec{w})$ such that $\sigma \models \text{fresh}(r; \vec{w})$ if and only if r does not occur in $\vec{w}\sigma$.

In this Section we give our axioms A . We use $\theta(\vec{v})$ as a shorthand for $\text{true} \parallel \theta(\vec{v})$ (i.e. axioms which are satisfied in any assignment of the free variables in \vec{v}).

A. Basic Properties

In Figure 4 we give basic axioms, mostly proposed in [BC14]. These state that \sim forms an equivalence relation on sets of terms, together with some essential properties of equality and conditional branching. In particular, we make regular use of the axiom $\vec{v}, [\text{EQ}(x, y)]u(x) : z \sim \vec{v}, [\text{EQ}(x, y)]u(y) : z$. The soundness proof of the basic axioms not already present in [BC14] can be found in Appendix C. We remark that Axiom 9 can be immediately derived from axiom IFEVAL in [BC16].

$z_1, [\text{true}]x : y, z_2 \sim z_1, x, z_2$	(1)
$z_1, [\text{false}]x : y, z_2 \sim z_1, y, z_2$	(2)
for $f : s_1 \times \dots \times s_n \rightarrow s$ $\vec{v}_1, x_1, \dots, x_n \sim \vec{v}_2, y_1, \dots, y_n$ $\rightarrow \vec{v}_1, f(x_1, \dots, x_n) \sim \vec{v}_2, f(y_1, \dots, y_n)$	(3)
$\vec{x} \sim \vec{x}$	(4)
$\vec{x} \sim \vec{y} \rightarrow \vec{y} \sim \vec{x}$	(5)
$\vec{x} \sim \vec{y} \wedge \vec{y} \sim \vec{z} \rightarrow \vec{x} \sim \vec{z}$	(6)
If p projects and permutes onto a sublist, $\vec{x} \sim \vec{y} \rightarrow p(\vec{x}) \sim p(\vec{y})$	(7)
$\text{EQ}(x, x) \sim \text{true}$	(8)
$\vec{v}, [\text{EQ}(x, y)]u(x) : w \sim \vec{v}, [\text{EQ}(x, y)]u(y) : w$	(9)
For $c \in \{\text{true}, \text{false}, \text{fail}\} \cup \mathcal{D} \cup \mathcal{H}$ $x \sim c \leftrightarrow \vec{v}, x \sim \vec{v}, c$	(10)

Figure 4. Basic axioms

We now state an important lemma that allow us to talk about semantic equality of terms.

Lemma 9. (*Rewriting With Equalities*) $\text{EQ}(x, y) \sim \text{true}$ if and only if $\vec{v}, x \sim \vec{v}, y$ for any sequence of terms \vec{v} . That is,

$\text{EQ}(x, y) \sim \text{true}$ if and only if x and y are indistinguishable in any context.

This Lemma follows from the basic axioms. The complete proof can be found in Appendix C. We remark that this Lemma, seen here as a consequence of the axioms, also follows immediately from axiom EQCONG in [BC16].

Remark 10. Since $\text{EQ}(x, y) \sim \text{true}$ means we can replace x by y everywhere, we use the notation $x = y$ in this case. This considerably simplifies long formulae, but one must note the distinction between *syntactic* equality $x \equiv y$, which means that x and y are the same terms from \mathcal{T} , and *semantic* equality $x = y$, which means that x and y are (possibly) different terms representing, *up to a negligible probability*, the same underlying object. Obviously by the computational interpretation of EQ, if $x \equiv y$ then $x = y$.

Function symbols applied to terms of sort bool in the symbolic model behave exactly as the corresponding operations on the booleans $\{0, 1\}$ in the computational model. We state the corresponding axioms in Figure 5 and prove them sound in Appendix C. Note that the boolean axioms described here could be easily derived from the ones in [BC16], given a proper axiomatisation of the logical connectives.

As a consequence of these axioms, we can manipulate terms of sort bool within sets of terms in an intuitive way. We use this manipulation regularly, without citing particular axioms, in the proof of the symbolic theorem in Section VI.

$b = [b]\text{true} : \text{false}$	(11)
$x = [b]x : x$	(12)
$[b_1]([b_2]u : v) : w = [b_1]([b_1 \rightarrow b_2]u : v) : w$	(13)
$[b_1]u : ([b_2]v : w) = [b_1]u : ([-b_1 \rightarrow b_2]v : w)$	(14)
<ul style="list-style-type: none"> • For any $f : s_1 \times \dots \times s_n \rightarrow s$ and for any $\vec{v} \equiv v_1, \dots, v_n$ and $\vec{w} \equiv w_1, \dots, w_n$ with each $v_i, w_i \in s_i$, 	
$f([b]\vec{v} : \vec{w}) = [b]f(\vec{v}) : f(\vec{w})$	(15)
<ul style="list-style-type: none"> • Let $f_1(X_1, \dots, X_n), f_2(Y_1, \dots, Y_m)$ be <i>boolean functions</i> (propositional terms built using the connectives $\{\neg, \wedge, \vee, \rightarrow\}$ and with n and m propositional variables, respectively). If $f_1(X_1, \dots, X_n) \Leftrightarrow f_2(Y_1, \dots, Y_m)$, 	
$f_1(b_1, \dots, b_n) = f_2(b_1, \dots, b_m)$	(16)

Figure 5. Boolean axioms

We state the following as an example of the kind of formal deduction that is possible:

Example 11. The formula $b_1 \rightarrow b_2 = [b_1]b_2 : \text{true}$ is true in any model for the axioms previously listed.

Next we present an axiom that is non-trivial and has not appeared in the literature before. The new axiom is used to prove a formula $\vec{v}_1, u_1 \sim \vec{v}_2, u_2$ by splitting it into multiple cases. In order for this strategy to work, the cases must form, with overwhelming probability, a perfect partition of the space of possibilities. This is formalised in the following definition:

Definition 12. Let I be a finite indexing set such that $(b^i)_{i \in I}$ is a family of terms of sort bool. If the formula

$$\bigvee_{i \in I} b^i \wedge \bigwedge_{i, j \in I, i \neq j} \neg(b^i \wedge b^j) = \text{true}$$

holds, then we say $(b^i)_{i \in I}$ is a *partition*.

We remark that if $(b^i)_{i \in I}$ is a partition, then, with overwhelming probability, there is exactly one $i \in I$ such that $[[b^i]]_{\eta, \rho}^\sigma = 1$.

If $(b_1^i)_{i \in I}$ and $(b_2^i)_{i \in I}$ are partitions (with the same indexing set), then we define the *case disjunction axiom*:

$\bigwedge_{i \in I} (\vec{v}_1, b_1^i, [b_1^i]u_1 \sim \vec{v}_2, b_2^i, [b_2^i]u_2) \rightarrow \vec{v}_1, u_1 \sim \vec{v}_2, u_2$	(17)
---	------

As we prove in Appendix C, this axiom is sound in all computational models.

B. Cryptographic Axioms

The core of the proof relies on axioms representing our *cryptographic* assumptions. To our knowledge, these axioms have not appeared in this form before, but the soundness proofs closely resemble others in the literature.

In Section V-C, we present some logical consequences of these axioms that are what we actually use in the symbolic proof.

The randomised version of the strong secrecy axiom states that if an encryption key k appears in a sequence of terms in essentially the way described in the AE-AD Privacy game - only as an encryption key and never with the same nonce for different plaintexts (and associated data) - then every encryption under k can be replaced with a random string (derived from a term of sort longnonce) while preserving indistinguishability. We formalise the constraint on k as follows:

Let the constraint $\text{enckey}(k; \vec{v})$ be true if and only if the following conditions are satisfied:

- In \vec{v} , the term k only appears in the position of an encryption key with a term of sort nonce in the position of the nonce and a term of sort attribute in associated data position,
- If two encryptions under k appearing in \vec{v} share the same term in the position of the nonce, then the terms in the plaintext and associated data positions are also identical.

Given a sequence of terms \vec{v} , let $\$_k(\vec{v})$ be the sequence of terms obtained by replacing every encryption $\text{enc}(m, a, r, k)$ under k appearing in \vec{v} with a term $\$(\text{enc}(m, a, r, k), r')$, where r' is a fresh term of sort longnonce that is indexed by the encryption term. Indexing the long nonces this way means that, if the same encryption term appears more than once in \vec{v} , then it is replaced with the same random string in $\$_k(\vec{v})$.

When encryption is randomised, using fresh nonces in the encryption terms ensures that ciphertexts created at different times will not collide (up to a negligible probability). So encryptions created at different times can be replaced by different random strings, regardless of the underlying plaintext,

without an adversary being able to distinguish the change. We do not have this guarantee when encryption is deterministic: encryptions of the same message (i.e. wraps of the same key) created at different times will give the same ciphertexts and must be replaced by the same random string. So for the deterministic version of our strong secrecy axiom, instead of replacing all encryptions under k at once we only replace the encryptions of a particular term. In addition we need to check in the premise of the axiom that the computational interpretation of this term is different to the interpretations of other (syntactically different) terms encrypted under k . For example, the encryption of $\pi_1(x, y)$ will need to be replaced by the same random string as the encryption of x , even though the plaintexts are syntactically different.

Let us define $\mathbb{S}_k^d(\vec{v}, m, a)$ as the sequence of terms obtained by replacing every encryption $\text{enc}(m, a, r, k)$ under k appearing in \vec{v} with a term $\mathbb{S}(\text{enc}(m, a, r, k), r')$, where r' is a fresh term of sort longnonce that is indexed by (m, a, k) . Note that r' must not depend on r as, in the deterministic interpretation of encryption, r is ignored.

<ul style="list-style-type: none"> • <i>Strong Secrecy of Wrapping (randomised):</i> 	$\text{enckey}(k; \vec{v}) \parallel \vec{v} \sim \mathbb{S}_k(\vec{v}) \quad (18)$
<ul style="list-style-type: none"> • <i>Strong Secrecy of Wrapping (deterministic):</i> 	$\text{enckey}(k; \vec{v}) \parallel \left(\bigwedge_{\substack{\{m'\}_{k,a}^r \in \text{st}(\vec{v}) \\ m' \neq m}} (\text{EQ}(m, m')) = \text{false} \right) \rightarrow \vec{v} \sim \mathbb{S}_k^d(\vec{v}, m, a) \quad (19)$
<ul style="list-style-type: none"> • <i>Integrity of Wrapping:</i> 	$\text{enckey}(k; x) \parallel \text{wf}(\text{uwp}(x, k)) = \left(\text{wf}(x) \wedge \bigvee_{\{m\}_{k,a}^r \in \text{st}(x)} \text{EQ}(x, \{m\}_{k,a}^r) \right) \quad (20)$
<ul style="list-style-type: none"> • <i>Correctness of Wrapping:</i> 	$\text{uwp}(\{m\}_{k,a}^r, k) = m \quad (21)$
<ul style="list-style-type: none"> • <i>Correctness of Projections:</i> 	$\pi_1(x, y, z) = x \quad (22)$
	$\pi_2(x, y, z) = y \quad (23)$
	$\pi_3(x, y, z) = z \quad (24)$
<ul style="list-style-type: none"> • <i>Random String Length Consistency: If $K \subseteq \mathcal{K}$</i> 	$\mathbb{S}_k(\mathfrak{o}_{k,A}(\vec{v})) \sim \mathbb{S}_k(\vec{v}) \quad (25)$
<ul style="list-style-type: none"> • <i>Failure Propagation:</i> 	$\{\text{fail}\}_{x,y}^z = \text{fail} = \text{uwp}(\text{fail}, x) \quad (26)$

Figure 6. Cryptographic axioms

Provided that the encryption and decryption symbols are together interpreted as a secure deterministic (resp. randomised) AE-AD scheme with some simple correctness assumptions,

the deterministic (resp. randomised) version of these axioms are sound. We prove this in Appendix C.

C. Consequences of the Axioms

While assuming AE-AD security gives us that ciphertexts are indistinguishable from random strings, we actually require a significantly weaker result in order to prove the security of key wrapping APIs: namely, that the encryptions of keys are indistinguishable from *encryptions* of fresh random strings (of the correct length). We call this the weak secrecy axiom.

For this axiom we have to introduce notation for replacing plaintexts inside encryptions. For a vector of terms \vec{v} , we denote by $\mathfrak{o}_{k,A}(\vec{v})$ the result of replacing any wrap $\{k'\}_{k,a}^r \in \text{st}(\vec{v})$ with $a \in A$ and $k \in \mathcal{K}$, by $\{r_{k,k'}\}_{k,a}^r$ where $r_{k,k'}$ denotes a fresh key (indexed by (k, k') in the same way as for \mathbb{S}_k and \mathbb{S}_k^d). For example, if $\vec{v} \equiv \pi_1(\{k'\}_{k,a}^r, \{k''\}_{k,a'}^{r'}, \{k\}_{k',a}^r)$, then $\mathfrak{o}_{k,\{a,a'\}}(\vec{v}) \equiv \pi_1(\{r_{k,k'}\}_{k,a}^r, \{r_{k,k''}\}_{k,a'}^{r'}, \{k\}_{k',a}^r)$.

So that the weak secrecy axiom holds for both randomised and deterministic encryption, we only allow this substitution for terms where the plaintexts are in fact constants of sort key (as opposed to more general terms whose computational interpretations are equal to those of keys). This is the wellused constraint, formally defined below.

We also prove that the weak secrecy axiom and the integrity axiom, which hold when a key only occurs in the encryption key position in a certain set of terms, also hold in a larger class of terms where the key is *usable*: as well as encryption key positions, the key might also occur in the plaintext position of an encryption, provided that the encryption key is itself usable. This is the case we will mostly encounter in our proofs. The formal definition of the usablekey constraint is given below.

<ul style="list-style-type: none"> • <i>Weak Secrecy of Wrapping:</i> 	$\text{enckey}(k; \vec{v}), \text{wellused}(a; \vec{v}; k) \parallel \vec{v} \sim \mathfrak{o}_{k,a}(\vec{v}) \quad (27)$
<ul style="list-style-type: none"> • <i>Usable Secrecy of Wrapping:</i> 	$\text{usablekey}(k; \vec{v}), \text{wellused}(a; \vec{v}; k) \parallel \vec{v} \sim \mathfrak{o}_{k,a}(\vec{v}) \quad (28)$
<ul style="list-style-type: none"> • <i>Usable Integrity of Wrap:</i> 	$\text{usablekey}(k; x) \parallel \text{wf}(\text{uwp}(x, k)) = \text{wf}(x) \wedge \bigvee_{\{m\}_{k,a}^r \in \text{st}(x)} \text{EQ}(x, \{m\}_{k,a}^r) \quad (29)$
<ul style="list-style-type: none"> • <i>Key Freshness:</i> 	$\text{usablekey}(k, t) \parallel \text{EQ}(k, t) = \text{false} \quad (30)$

Figure 7. Consequences of the cryptographic axioms

Definition 13. Let X be a set of terms. For all $k \in \mathcal{K}$, we say k is *usable* in X – and write $\text{usablekey}(k; X)$ – if the following holds:

- 1) Either $\text{enckey}(k; X)$,
- 2) or if k appears in a plaintext position $p.a$ in X , then $X|_p \equiv \{k\}_{k',a}^r$ such that

- a) $k' \in \mathcal{K}$ is usable in X , a is well used for k' in X and r is a term of sort nonce
- b) if r occurs in a term $\{k''\}_{k',a'}^r \in \text{st}(X)$, then $k'' \equiv k$ and $a' \equiv a$.

Definition 14. Let X be a set of terms. We say that the attribute a is *well-used* for the key k in X (written $\text{wellused}(a; X; k)$) if, for every $\{m\}_{k,a}^r \in \text{st}(X)$, m is a constant of sort key.

The proof that these axioms are consequences of the previous cryptographic axioms is proved in Appendix C.

D. Policy Axioms

We conclude this Section by giving, in Figure 8, axioms reflecting the validity of a policy. These axioms are sound under the conditions of Section III.

As before, there are three criteria: the first is that the policy respects the sacrifice function, in the sense that if the policy accepts the pair (x, a) , then x must be semantically equal to an attribute b such that $b \in \text{sacr}(a)$. The second criterion is that the policy never allows external keys to be used for wrapping, so the policy will never accept (x, a) if $a \in \mathcal{E}$. Finally, the policy must forbid the creation of key cycles by enforcing a strict partial ordering on attributes. Additionally we give two useful consequences of these axioms, namely stating that compromise is monotonic in the API execution and that the compromise function respects the policy. These axioms are sound as soon as the function implementing the policy is valid with respect to **sacr**.

Policy axioms:	
for $a \in \mathcal{D}$: $\text{P}(x, a) = \left(\text{P}(x, a) \wedge \bigvee_{b \in \text{sacr}(a)} \text{EQ}(x, b) \right)$	(31)
for $a \in \mathcal{E}$: $\text{P}(x, a) = \text{false}$	(32)
• There exists a strict partial ordering \prec on \mathcal{D} such that:	
$(\text{P}(a, b) = \text{true}) \rightarrow a \prec b$	(33)
Consequences of the policy axioms and the definition of comp :	
• If there is a sequence of actions inducing a transition from state q to state q' : $\text{comp}_q(x) \rightarrow \text{comp}_{q'}(x) = \text{true}$	
	(34)
• For $x, y \in \mathcal{T} \setminus \mathcal{H}$:	
$(\text{comp}_q(y) \wedge \text{P}(x, y) \rightarrow \text{comp}_q(x)) = \text{true}$	(35)

Figure 8. Policy axioms

VI. PROOF OF THE SYMBOLIC THEOREM

A detailed proof is provided in Appendix D. We provide here an outline of the proof.

Before tackling the main Theorem, we prove that the frames ϕ^* and ϕ^\dagger , which are *a priori* very complicated sequences of terms, can be rewritten in a relatively simple form. Let K_0 be the set of honestly-generated keys initially stored by the API.

We start by identifying sets of terms in which keys in K_0 appear only as expected: if they are in a plaintext position

then they are protected by a properly-constructed wrap or known to have a compromised attribute; if they appear in an unwrap position (which cannot be simplified using the integrity and correctness axioms) then they are known to have a compromised attribute; and if they appear in a wrapping position then the key's attribute is internal and, either the wrap is of a key from K_0 and its corresponding attribute, or the associated data of the wrap is a compromised attribute. We remark that – by the policy axioms – uncompromised keys are usable in such sets.

We then identify a sequence $(F^i)_{i \in I}$ of boolean terms, in which keys in K_0 appear as above, that partition the space of possibilities and give sufficient information to determine whether or not any given handle or attribute is compromised. We prove a number of invariants showing that keys in K_0 only appear as expected in $[F^i]\phi$ and $[F^i]\text{val}_q(h)$ and so uncompromised keys are usable in these terms. Additionally, we show that under the conditions of the (F^i) , the internal state of the API is consistent with its initial configuration.

Since k^* is uncompromised, the usable secrecy axiom allows us to remove every instance of it from $[F^i]\phi$ until we reach a set of terms where replacing k^* by k^\dagger has no effect. The main theorem then follows by the case disjunction axiom (Axiom 17) applied to the partition $(F^i)_{i \in I}$.

VII. COMPOSING KEY WRAPPING WITH OTHER FUNCTIONALITIES

In this Section we demonstrate how our security property for a key wrapping API composes with the security of other functionalities offered by a cryptographic API, provided that keys used by the other functionalities are external. We will use encryption (of data, not keys) as an example but the argument can be easily adapted for other cryptographic primitives.

First, we define a game in which the adversary interacts with a key wrapping API and an encryption scheme that uses the external keys managed by the key wrapping API. For simplicity, here the algorithms `enc` and `dec` used by the encryption scheme (not the wrapping mechanism) take just two arguments, the first being the key. Our game is very similar to the one used to define API security in [KSW11].

The game is parameterised by a handle h^* pointing to the key k^* with external attribute a^* . The adversary has access to a real-or-random encryption oracle for k^* and a (genuine) decryption oracle for k^* . Much as in the real-or-random IND-CCA game, the adversary tries to guess whether or not the encryption oracle for k^* is real, without passing any of its outputs to the decryption oracle for k^* . But in addition, the adversary is allowed to wrap and unwrap keys (according to the policy from the key wrapping API), encrypt and decrypt as normal under any external key other than k^* . Finally the adversary may also corrupt any key, provided that they do not compromise h^* . We formalise this below.

Definition 15. Let $\text{API} = (\mathcal{K}, \mathcal{D}, \mathcal{E}, \mathcal{H}, \text{wm}, \text{P})$ be a key wrapping API and let $\Pi = (\text{keygen}, \text{enc}, \text{dec})$ be an encryption scheme. Then we say $\text{API}' = (\text{API}, \Pi)$ is a secure

encryption and wrapping API if, for all integers m and n , all $\vec{a} = a_1, \dots, a_n \in \mathcal{D}^n$ with $a_1 = a^* \in \mathcal{E}$ and all polynomial-time adversaries \mathcal{A} , the following advantage is a negligible function of η :

$$\text{Adv}_{\text{API}'}^{\text{IND-CCA}}(\mathcal{A}) := \mathbb{P}_b [\text{Exp}_b^{\text{COMP}}(\mathcal{A}, \vec{a}, m) = b] - \frac{1}{2}$$

where the experiments $\text{Exp}_b^{\text{COMP}}(\mathcal{A}, \vec{a}, m)$ for $b \in \{0, 1\}$ and the new oracles are given in Figures 9 and 10.

```

Experiment  $\text{Exp}_b^{\text{COMP}}(\mathcal{A}, a_1, \dots, a_n, m)$ :
st.C  $\leftarrow 0$ 
st.V  $\leftarrow \top$ 
st.cor  $\leftarrow []$ 
st.e  $\leftarrow []$ 
For  $1 \leq i \leq n$ 
   $k_i \leftarrow \text{keygen}(1^\eta)$ 
  st.val( $h_i$ )  $\leftarrow (k_i, a_i)$ 
st.H  $\leftarrow \{h_1, \dots, h_n\}$ 
( $h^*, a^*, k^*$ )  $\leftarrow (h_1, a_1, k_1)$ 
st.T  $\leftarrow a^*$ 
 $\mathcal{O}_{\text{API}} \leftarrow (\mathcal{O}^{\text{wrap}}, \mathcal{O}^{\text{unwrap}}, \mathcal{O}^{\text{corrupt}}, \mathcal{O}_b^{\text{enc}}, \mathcal{O}_b^{\text{dec}})$ 
 $b' \leftarrow \mathcal{A}^{\mathcal{O}_{\text{API}}}$ 
If st.V =  $\top$  Return  $b'$ 
 $b'' \leftarrow \{0, 1\}$ 
Return  $b''$ 

```

Figure 9. Cryptographic API Security Experiment

Oracle $\mathcal{O}_b^{\text{enc}}(h, m)$:	Oracle $\mathcal{O}_b^{\text{dec}}(h, c)$:
st.C \leftarrow st.C + 1	st.C \leftarrow st.C + 1
If st.C > m Return \perp	If st.C > m Return \perp
(k, a) \leftarrow st.val(h)	If $c \in$ st.e then st.V $\leftarrow \perp$
If $a \notin \mathcal{E}$ Return \perp	(k, a) \leftarrow st.val(h)
If $h = h^*$	If $a \notin \mathcal{E}$ Return \perp
If b	Return dec(k, c)
$c \leftarrow$ enc(k, m)	
Else	
$m' \leftarrow \{0, 1\}^{ m }$	
$c \leftarrow$ enc(k, m')	
st.e \leftarrow st.e $\cup \{c\}$	
Return c	
Else	
Return enc(k, m)	

Figure 10. Encryption and Decryption Oracles

We obtain the following composition theorem (the proof is in Appendix E).

Theorem 3. *Let API be key wrapping API satisfying the hypotheses of Theorem 1 and let Π be an IND-CCA secure encryption scheme. Then $\text{API}' = (\text{API}, \Pi)$ is a secure encryption and wrapping API.*

VIII. APPLICATIONS

A. Refinement of PKCS#11

We view PKCS#11 as the composition of a key wrapping API and an encryption scheme in the sense of Theorem 3. We describe a refinement of Version 2.4 of the PKCS#11 standard, such that this composition is secure in the sense of Theorem 3. Moreover, our refinement is less restrictive and our security guarantees are stronger than in [Kün15].

The PKCS#11 standard does not make a distinction between the key management module of a cryptographic token and its other functionalities [OAS15a]. However, as the wrap and decrypt attack from [BCFS10] demonstrates, not separating keys for different functionalities can lead to attacks. Therefore to obtain a secure refinement of PKCS#11 we first forbid the creation of keys with attributes that can be used for multiple roles and we do not allow attributes to be modified. We remark that the standard explicitly supports the use of additional restrictions on attributes such as these.

Our final additional assumption is that the wrapping mechanism is a secure AE-AD encryption scheme (for example AES-GCM, supported in PKCS#11 since version 2.4 of the standard) that binds attributes properly when wrapping. Note that all our refinements are explicitly permitted by the PKCS#11 specification.

Then, Theorem 3 shows that an IND-CCA secure encryption mechanism using external keys can be securely combined with the key management functionalities of PKCS#11. More details of this analysis can be found in Appendix F.

B. The Kremer-Steel-Warinschi API

In this Section we show how our result can be used to encapsulate security proofs of other computational APIs that have appeared in the literature. As an example we view the API in [KSW11] as an instance of a generic API described by our model and verify its security.

The authors of [KSW11] give a security definition for APIs with wrapping and encryption mechanisms. As in our composition result, the adversary is not supposed to distinguish between real and fake encryptions under a particular challenge key. The authors go on to describe an example implementation, in which the wrapping policy is a simple hierarchy on keys with external keys at the bottom level, and prove its security. Clearly this is a valid policy according to our definition. Therefore Theorem 3 subsumes the security proof given for this particular API design in [KSW11].

IX. CONCLUSION

We give here a general definition of a key wrapping API, parameterised by a wrapping mechanism and a wrapping policy. We provide a set of simple conditions for a key wrapping policy to be valid. Namely, forbidding creation of wrapping cycles, not allowing wraps under external keys and respecting the enterprise-level policy. We prove that if, in addition to these policy conditions, the key wrapping mechanism is implemented by an AE-AD encryption scheme and attributes are properly transmitted, the external keys are indistinguishable from random values and the key wrapping API can be securely composed with an encryption scheme.

Our strong secrecy notion for external keys, together with our generic notion of valid policies, allows us to prove that a secure key wrapping API may be securely composed with an encryption scheme. We are then able to give a configuration of the key wrapping functionality of PKCS#11 such that the

encryption mechanism is as secure as when fresh keys are used.

Since our proof technique relies on the CCSA from [BC14], our theorem is derived from a set of relatively minimal axioms representing computational assumptions. These axioms are generic enough to be proven sound under a number of different cryptographic hypotheses, notably deterministic and randomised variants of authenticated encryption.

Ultimately, the hypotheses required for our security proof - particularly the need to securely bind the attributes of keys when wrapping - may be far from what currently happens in practice. Nevertheless, we argue that this assumption is essential for security, since one must enforce a separation between keys used for wrapping and keys used for other primitives. What we offer is a clear blueprint for building secure key management APIs in the future:

- 1) Explicitly set out which keys are intended for wrapping (internal) and which are intended for other uses (external) and the intended security relationships between keys (a sacrifice function).
- 2) Build an API that uses a secure AE-AD encryption scheme to bind the attributes of keys to the wraps and where, when wrapping keys, there is an *explicit* check performed on attributes (an API security policy).
- 3) Verify that the API security policy is acyclic, respects the sacrifice function and forbids the use of external keys for wrapping.

In this way, one guarantees that external keys are as secure as possible.

X. ACKNOWLEDGEMENTS

This work was partially supported by the European Union's 7th Framework Program (FP7/2007-2013) under grant agreement n. 609611 (PRACTICE). The authors would like to thank Bogdan Warinschi, Martijn Stam and Karthikeyan Bhargavan for their useful feedback on the paper. The authors would also like to thank the anonymous reviewers for their valuable comments.

REFERENCES

- [AFL13] Pedro Adão, Riccardo Focardi, and Flaminia L. Luccio. Type-based analysis of generic key management apis. In *2013 IEEE 26th Computer Security Foundations Symposium, New Orleans, LA, USA, June 26-28, 2013*, pages 97–111. IEEE, 2013.
- [AR07] Martín Abadi and Phillip Rogaway. Reconciling two views of cryptography (the computational soundness of formal encryption). *J. Cryptology*, 20(3):395, 2007.
- [BC14] Gergei Bana and Hubert Comon-Lundh. A computationally complete symbolic attacker for equivalence properties. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pages 609–620, 2014.
- [BC16] Gergei Bana and Rohit Chadha. Verification methods for the computationally complete symbolic attacker based on indistinguishability. Cryptology ePrint Archive, Report 2016/069, 2016. <http://eprint.iacr.org/>.
- [BCFS10] Matteo Bortolozzo, Matteo Centenaro, Riccardo Focardi, and Graham Steel. Attacking and fixing PKCS#11 security tokens. In *Proceedings of the 17th ACM Conference on Computer and Communications Security, CCS 2010, Chicago, Illinois, USA, October 4-8, 2010*, pages 260–269, 2010.
- [BFBK⁺12] Romain Bardou, Riccardo Focardi, Yusuke Kawamoto, Lorenzo Simionato, Graham Steel, and Joe-Kai Tsay. Efficient padding oracle attacks on cryptographic hardware. In *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 608–625. Springer, 2012.
- [CC09] Christian Cachin and Nishanth Chandran. A secure cryptographic token interface. In *Proceedings of the 22nd IEEE Computer Security Foundations Symposium, CSF 2009, Port Jefferson, New York, USA, July 8-10, 2009*, pages 141–153, 2009.
- [CFL13] Matteo Centenaro, Riccardo Focardi, and Flaminia L. Luccio. Type-based analysis of key management in pkcs#11 cryptographic devices. *Journal of Computer Security*, 21(6):971–1007, 2013.
- [CGH12] David Cash, Matthew Green, and Susan Hohenberger. New definitions and separations for circular security. In *Public Key Cryptography - PKC 2012 - 15th International Conference on Practice and Theory in Public Key Cryptography, Darmstadt, Germany, May 21-23, 2012. Proceedings*, volume 7293 of *Lecture Notes in Computer Science*, pages 540–557. Springer, 2012.
- [Clu03] Jolyon Clulow. On the security of PKCS#11. In *Cryptographic Hardware and Embedded Systems - CHES 2003, 5th International Workshop, Cologne, Germany, September 8-10, 2003. Proceedings*, pages 411–425, 2003.
- [CS09] Véronique Cortier and Graham Steel. A generic security API for symmetric key management on cryptographic devices. In *Computer Security - ESORICS 2009, 14th European Symposium on Research in Computer Security, Saint-Malo, France, September 21-23, 2009. Proceedings*, pages 605–620, 2009.
- [CSW12] Véronique Cortier, Graham Steel, and Cyrille Wiedling. Revoke and let live: a secure key revocation API for cryptographic devices. In *the ACM Conference on Computer and Communications Security, CCS'12, Raleigh, NC, USA, October 16-18, 2012*, pages 918–928, 2012.
- [DKS08] Stéphanie Delaune, Steve Kremer, and Graham Steel. Formal analysis of PKCS#11. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium, CSF 2008, Pittsburgh, Pennsylvania, 23-25 June 2008*, pages 331–344, 2008.
- [DKS10] Stéphanie Delaune, Steve Kremer, and Graham Steel. Formal security analysis of PKCS#11 and proprietary extensions. *Journal of Computer Security*, 18(6):1211–1245, 2010.
- [KKS13] Steve Kremer, Robert Künnemann, and Graham Steel. Universally composable key-management. In *Computer Security - ESORICS 2013 - 18th European Symposium on Research in Computer Security, Egham, UK, September 9-13, 2013. Proceedings*, pages 327–344, 2013.
- [KSW11] Steve Kremer, Graham Steel, and Bogdan Warinschi. Security for key management interfaces. In *Proceedings of the 24th IEEE Computer Security Foundations Symposium, CSF 2011, Cernay-la-Ville, France, 27-29 June, 2011*, pages 266–280, 2011.
- [Kün15] Robert Künnemann. Automated backward analysis of pkcs#11 v2.20. In *Principles of Security and Trust - 4th International Conference, POST 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*, volume 9036 of *Lecture Notes in Computer Science*, pages 219–238. Springer, 2015.
- [NIS12] NIST Special Publication 800-57: Recommendation for Key Management - Part 1: General (Revision 3), July 2012.
- [OAS15a] PKCS#11 cryptographic token interface base specification version 2.40, April 2015. Latest version: <http://docs.oasis-open.org/pkcs11/pkcs11-base/v2.40/pkcs11-base-v2.40.html>.
- [OAS15b] PKCS#11 cryptographic token interface current mechanisms specification version 2.40, April 2015. Latest version: <http://docs.oasis-open.org/pkcs11/pkcs11-curr/v2.40/pkcs11-curr-v2.40.html>.
- [Rog02] Phillip Rogaway. Authenticated-encryption with associated-data. In *Proceedings of the 9th ACM Conference on Computer and Communications Security, CCS 2002, Washington, DC, USA, November 18-22, 2002*, pages 98–107, 2002.
- [RS06] Phillip Rogaway and Thomas Shrimpton. Deterministic authenticated encryption: A provable-security treatment of the key-wrap problem. In *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications*

APPENDIX A COMPUTATIONAL INTERPRETATION

Any term t is interpreted as a PPT \mathcal{M}_t with two random tapes and one input tape. The interpretation of the term t , for a security parameter η , honest randomness ρ_1 and adversarial randomness ρ_2 is the bitstring $s = \mathcal{M}_t(1^\eta; \rho_1; \rho_2)$, which denotes \mathcal{M}_t computing on input 1^η with random tapes ρ_1 and ρ_2 . If t has sort `bool` then $s \in \{0, 1\}$. The honest randomness is used for fixing the interpretation of nonces and keys whereas the adversarial randomness is used in the interpretation of the adversarial functions $g \in \mathcal{G}$. For $\rho = (\rho_1, \rho_2)$, we write $\llbracket t \rrbracket_{\eta, \rho} = \mathcal{M}_t(1^\eta; \rho_1; \rho_2)$. We define \mathcal{M}_t inductively as follows:

For every n -ary function $f \in \mathcal{F}$ we fix a (deterministic) Turing machine $\llbracket f \rrbracket$ taking $n + 1$ arguments (the additional argument is the security parameter). In order for our model to be a computational model, we insist that these machines should satisfy some conditions.

- $\mathcal{M}_{\text{true}}$ takes no inputs and returns 1 (so $\llbracket \text{true} \rrbracket \in D_{\text{bool}}$)
- $\mathcal{M}_{\text{false}}$ takes no inputs and returns 0 (so $\llbracket \text{false} \rrbracket \in D_{\text{bool}}$)
- $\mathcal{M}_{\text{fail}}$ takes no inputs and returns a fixed bitstring (to denote failure), which we label ζ

$$\mathcal{M}_{\text{if_then_else_}}(p, s, s_1, s_2) = \begin{cases} s_1, & \text{if } s = 1, \\ s_2, & \text{if } s = 0. \end{cases}$$

$$\mathcal{M}_{\text{_and_}}(p, s_1, s_2) = \begin{cases} 1, & \text{if } s_1 = s_2 = 1, \\ 0, & \text{otherwise.} \end{cases}$$

$$\mathcal{M}_{\text{_or_}}(p, s_1, s_2) = \begin{cases} 1, & \text{if } s_1 = 1 \text{ or } s_2 = 1, \\ 0, & \text{otherwise.} \end{cases}$$

$$\mathcal{M}_{\text{_not_}}(p, s) = \begin{cases} 1, & \text{if } s = 0, \\ 0, & \text{otherwise.} \end{cases}$$

$$\mathcal{M}_{\text{EQ}}(p, s_1, s_2) = \begin{cases} 1, & \text{if } s_1 = s_2, \\ 0, & \text{otherwise.} \end{cases}$$

$$\mathcal{M}_{\text{o}}(p) = 0^{\text{keylen}(1^\eta)}$$

$$\mathcal{M}_{\text{\$}}(p, s_1, s_2) = \text{the first } |s_1| \text{ bits of } s_2$$

We write `enc` and `dec` for the machines $\llbracket \text{enc} \rrbracket$ and $\llbracket \text{dec} \rrbracket$, respectively. We make the following assumptions about these algorithms:

- The machine `enc`, on input (p, s_m, s_a, s_r, s_k) , outputs the authenticated encryption of s_m under the key s_k with associated data s_a and nonce s_r . For deterministic encryption `enc` can just ignore s_r .

- The machine `dec`, on input (p, s_c, s_a, s_r, s_k) attempts to decrypt s_c under the key s_k with associated data s_a and nonce s_r . If the decryption fails, the machine outputs ζ .

The n -ary adversarial function g_i is interpreted as an arbitrary PPT taking $n+1$ arguments. When used for interpreting a term, it is given randomness ρ_2 . In this way all the adversarial functions share the same randomness.

Every $n \in \mathcal{N}$ of sort `s` is interpreted as a PPT \mathcal{M}_s taking its randomness from ρ_1 according to its sort. In other words we fix a PPT for each sort, that given a random tape and the security parameter returns the value of a name of this sort. In order to ensure that all nonces and keys are sampled independently, we fix a (polynomial-time) extraction procedure such that, for any infinite sequence ρ_1 of bits, there is an infinite subsequence ρ_1^n of ρ_1 for each $n \in \mathcal{N}$. This extraction is such that the subsequences $\rho_1^{n_1}$ and $\rho_1^{n_2}$ are disjoint if $n_1 \neq n_2$. In this way, if each bit of ρ_1 is selected uniformly at random from $\{0, 1\}$, then the bits of ρ_1^n for each $n \in \mathcal{N}$ are independent and uniformly random elements of $\{0, 1\}$. The randomness used for the interpretation of a name n is ρ_1^n , hence every name is interpreted uniformly at random.

This interpretation of each sort is as expected, for example if k is of sort `key`, $\llbracket k \rrbracket_{\eta, \rho} = \text{keygen}(1^\eta; \rho_1^k)$ and this will always have length `keylen`(1^η). Further, we insist that the interpretations of names of sorts handle, attribute do not depend on ρ . Note that we separate the sorts nonce and longnonce. This is because a term r of sort `longnonce` will be used to replace a real ciphertext c by a random string of the same length (in particular, by $\$(c, r)$). Thus, the length of long nonces must be at least the length of arbitrary ciphertexts. On the other hand, ordinary nonces will be used by the encryption scheme and so it is unreasonable for them to be arbitrarily long.

We can now recall, from [BC14], the computational interpretation of an arbitrary term $t \in \mathcal{T}(\mathcal{F}, \mathcal{G}, \mathcal{N})$, given a security parameter η and a random sampling $\rho = \rho_1, \rho_2$:

- For $n \in \mathcal{N}$ of sort `s`, $\llbracket n \rrbracket_{\eta, \rho} = \mathcal{M}_s(1^\eta; \rho_1)$
- For $f \in \mathcal{F}$,

$$\llbracket f(t_1, \dots, t_n) \rrbracket_{\eta, \rho} = \llbracket f \rrbracket(1^\eta, \llbracket t_1 \rrbracket_{\eta, \rho}, \dots, \llbracket t_n \rrbracket_{\eta, \rho})$$

- For $g \in \mathcal{G}$

$$\llbracket g(t_1, \dots, t_n) \rrbracket_{\eta, \rho} = \llbracket g \rrbracket(1^\eta, \rho_2 \llbracket t_1 \rrbracket_{\eta, \rho}, \dots, \llbracket t_n \rrbracket_{\eta, \rho})$$

Note that the interpretation of names only depends on ρ_1 , whereas the interpretation of the adversarial functions only depends on ρ_2 (and the values of their arguments). In this way we separate the randomness used by the adversary to compute its successive queries and the randomness used by the API.

Finally, we fix the interpretation of the formulae $u_1, \dots, u_n \sim v_1, \dots, v_n$. As in [BC14], a formula of this form is interpreted, given a security parameter η and a sampling ρ , as the *computational indistinguishability* of the sequences of bitstrings $\llbracket u_1, \dots, u_n \rrbracket_{\eta, \rho}$ and $\llbracket v_1, \dots, v_n \rrbracket_{\eta, \rho}$. Formally, a computational model (an interpretation of the names, variables and function symbols) satisfies the formula $\psi \equiv u_1, \dots, u_n \sim$

v_1, \dots, v_n if and only if, for all polynomial-time adversaries \mathcal{A} with access to the random tape ρ_2 , the following advantage is negligible in η :

$$\text{Adv}_{\psi}^{\sim}(\mathcal{A}) := \mathbb{P} \left[\mathcal{A} \left(\llbracket u_1, \dots, u_n \rrbracket_{\eta, \rho}^{\sigma}; \rho_2 \right) = 1 \right] \\ - \mathbb{P} \left[\mathcal{A} \left(\llbracket v_1, \dots, v_n \rrbracket_{\eta, \rho}^{\sigma}; \rho_2 \right) = 1 \right],$$

where the probabilities are over the choice of ρ_1 and ρ_2 . Note that the distinguisher has access to the random tape used for computing the adversarial functions. In this way, we can consider that the interpretation of every function in \mathcal{G} , and the distinguisher, are different phases of the same adversary.

APPENDIX B REDUCING COMPUTATIONAL SECURITY

We define new computational security experiments that closely resemble the symbolic security properties proved in Theorem 2.

In the original computational security definitions, API executions have a variable structure, since queries can occur in any order. This is unlike the symbolic security properties where there are fixed sequences of output terms to compare. We therefore define new experiments where, in each experiment, the adversary executes a particular list of fully-specified actions.

Another important difference between the computational and symbolic security definitions is as follows: in the symbolic key secrecy property, an external key k^* is chosen *before* the API execution, and the adversary sees k^* in both ϕ^* and ϕ^\dagger , but the handle h^* may or may not point to k^* . So, to reflect this in our new computational games, we have one game for each possible test handle h_i and, at the start of the games, the adversary learns the value of a key that may or may not be pointed to by h_i .

Finally, in the symbolic security properties, we say that no adversary can distinguish between sets of terms that arise from an API execution. To model this, the adversary \mathcal{A} in the new experiments has two, independent components \mathcal{A}_1 and \mathcal{A}_2 : in the key secrecy experiments, \mathcal{A}_1 makes oracle queries determining an API execution, and \mathcal{A}_2 sees the inputs and outputs of these queries only if the API execution does not compromise the test handle and the test handle has an external attribute. Similarly, in the handle consistency experiments we have \mathcal{A}_1 who executes an API and \mathcal{A}_2 who tries to distinguish, from the bit 0, the value of a test for handle inconsistency.

Definition 16. Let $\mathbf{L} \in \text{Act}^m$, $1 \leq i \leq n$ and $b \in \{0, 1\}$. Then, we define the *restricted key secrecy experiments* $\text{Exp}_{\mathbf{L}, i, b}^{\text{KEYSEC}}$ and the *restricted handle consistency experiments* $\text{Exp}_{\mathbf{L}, i, b}^{\text{HDLCON}}$, where the formal definitions of these experiments and the relevant oracles are given in Figure 11.

We prove in Lemma 17 that security according to the experiments $\text{Exp}_b^{\text{KEYSEC}}$ and $\text{Exp}_b^{\text{HDLCON}}$ reduces to security according to the experiments $\text{Exp}_{\mathbf{L}, i, b}^{\text{KEYSEC}}$ and $\text{Exp}_{\mathbf{L}, i, b}^{\text{HDLCON}}$ for all $\mathbf{L} \in \text{Act}^m$ and $1 \leq i \leq n$. This is because the number of handles n and the number of fully-specified actions m are

assumed to be constant in the security parameter. Moreover, the latter experiments correspond exactly to the symbolic security properties proved in Theorem 2. This is how we relate our symbolic security properties to the computational definitions of security and obtain Theorem 1.

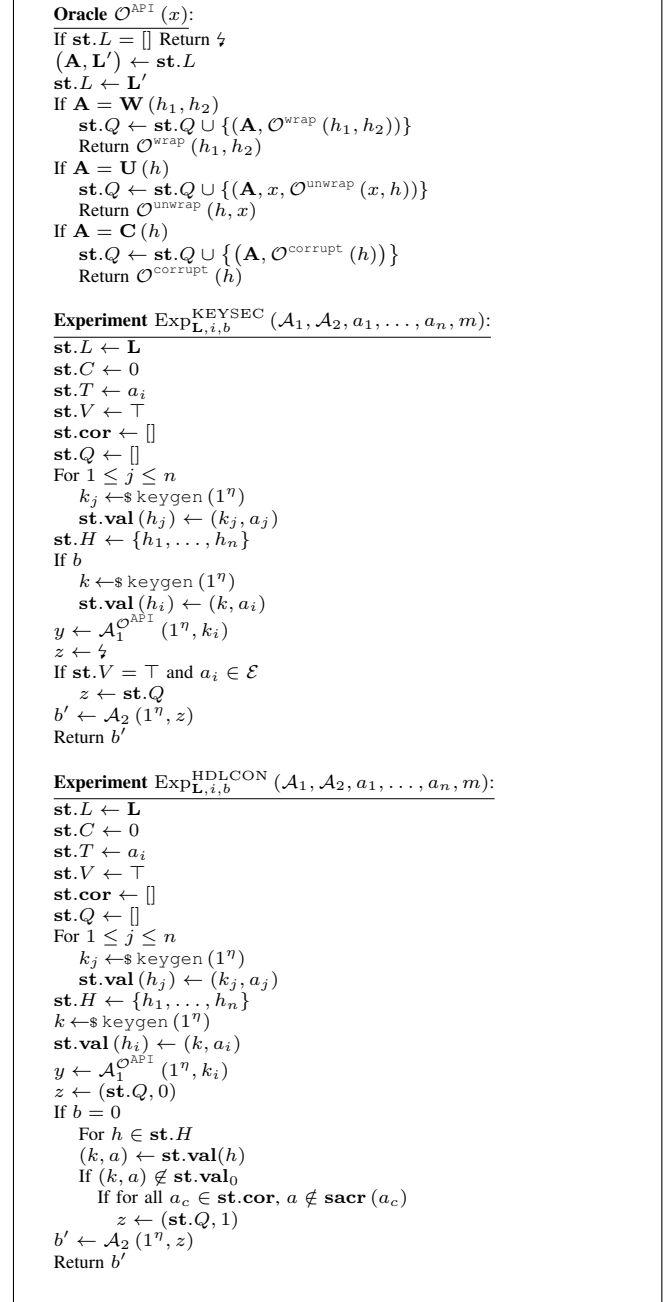


Figure 11. API oracle, Restricted Key Secrecy Experiment and Restricted Handle Consistency Experiment for (\mathbf{L}, i, b)

Lemma 17. *If, for all integers m and n , all $\vec{a} \in \mathcal{D}^n$, all polynomial-time adversaries \mathcal{A}_1 and \mathcal{A}_2 , all $\mathbf{L} \in \text{Act}$ and all*

$1 \leq i \leq n$, the advantages

$$\text{Adv}_{\text{API}^m, \mathbf{L}, i}^{\text{KEYSEC}}(\mathcal{A}_1, \mathcal{A}_2) := \mathbb{P}_b [\text{Exp}_{\mathbf{L}, i, b}^{\text{KEYSEC}}(\mathcal{A}_1, \mathcal{A}_2, \vec{a}, m) = b] - \frac{1}{2}$$

$$\text{Adv}_{\text{API}^m, \mathbf{L}, i}^{\text{HDLCON}}(\mathcal{A}_1, \mathcal{A}_2) := \mathbb{P}_b [\text{Exp}_{\mathbf{L}, i, b}^{\text{HDLCON}}(\mathcal{A}_1, \mathcal{A}_2, \vec{a}, m) = b] - \frac{1}{2}$$

are negligible in η , then API is a secure key wrapping API.

Proof. Consider an adversary \mathcal{A} in the original key secrecy experiment with m queries. For each $\mathbf{L} \in \text{Act}^m$ and $1 \leq i \leq n$ we construct an adversary $\mathcal{B}_{\mathbf{L}, i}$ for the same experiment who computes as follows: $\mathcal{B}_{\mathbf{L}, i}$ submits the handle h_i to the test oracle, then simulates \mathcal{A} , forwarding all queries and replies to its API oracle for as long as the queries from \mathcal{A} agree with \mathbf{L} . If \mathcal{A} tests h_i , $\mathcal{B}_{\mathbf{L}, i}$ sends the reply from its test oracle it received at the start of the game. If, at any point, \mathcal{A} deviates from \mathbf{L} or tests a handle other than h_i , then $\mathcal{B}_{\mathbf{L}, i}$ stops simulating \mathcal{A} , continues its execution of the API according to \mathbf{L} (submitting random inputs to the unwrap queries when necessary) and finally outputs a random bit. If the entire execution from \mathcal{A} agrees with \mathbf{L} and \mathcal{A} tests h_i , then $\mathcal{B}_{\mathbf{L}, i}$ returns the bit that \mathcal{A} returns.

For $b \in \{0, 1\}$, let

$$\mathcal{O}_b = (\mathcal{O}^{\text{wrap}}, \mathcal{O}^{\text{unwrap}}, \mathcal{O}^{\text{corrupt}}, \mathcal{O}_b^{\text{test}})$$

Then, for any adversary \mathcal{A} , let $\text{bad}(\mathcal{A})$ be the event that $\mathcal{A}^{\mathcal{O}_b}(1^\eta)$ compromises its test handle or its test handle does not have an external attribute. Without loss of generality this is independent of b . Then,

$$\mathbb{P} [\text{Exp}_b^{\text{KEYSEC}}(\mathcal{A}, a_1, \dots, a_n, m) = 1] = \frac{1}{2} \mathbb{P} [\text{bad}(\mathcal{A})] + \mathbb{P} [\mathcal{A}^{\mathcal{O}_b}(1^\eta) = 1 \cap \neg \text{bad}(\mathcal{A})]$$

so

$$\text{Adv}_{\text{API}^m}^{\text{KEYSEC}}(\mathcal{A}) = \mathbb{P} [\mathcal{A}^{\mathcal{O}_1}(1^\eta) = 1 \cap \neg \text{bad}(\mathcal{A})] - \mathbb{P} [\mathcal{A}^{\mathcal{O}_0}(1^\eta) = 1 \cap \neg \text{bad}(\mathcal{A})]$$

Let $E(\mathbf{L}, i)$ be the event that the entire execution from \mathcal{A} agrees with \mathbf{L} and \mathcal{A} tests h_i , which is independent of the bit b in the experiment. Without loss of generality, we assume that \mathcal{A} makes exactly m oracle queries apart from queries to the test oracle, and tests exactly one handle. Therefore $\sum_{\mathbf{L} \in \text{Act}^m} \sum_{i=1}^n \mathbb{P} [E(\mathbf{L}, i)] = 1$. Then we have

$$\begin{aligned} & \mathbb{P} [\mathcal{A}^{\mathcal{O}_b}(1^\eta) = 1 \cap \neg \text{bad}(\mathcal{A})] \\ &= \sum_{\mathbf{L} \in \text{Act}^m} \sum_{i=1}^n \mathbb{P} [\mathcal{A}^{\mathcal{O}_b}(1^\eta) = 1 \cap \neg \text{bad}(\mathcal{A}) \cap E(\mathbf{L}, i)] \\ &= \sum_{\mathbf{L} \in \text{Act}^m} \sum_{i=1}^n \mathbb{P} [\mathcal{B}_{\mathbf{L}, i}^{\mathcal{O}_b}(1^\eta) = 1 \cap \neg \text{bad}(\mathcal{B}_{\mathbf{L}, i}) \cap E(\mathbf{L}, i)] \end{aligned}$$

Now

$$\begin{aligned} & \mathbb{P} [\mathcal{B}_{\mathbf{L}, i}^{\mathcal{O}_b}(1^\eta) = 1 \cap \neg \text{bad}(\mathcal{B}_{\mathbf{L}, i}) \cap E(\mathbf{L}, i)] \\ &= \mathbb{P} [\mathcal{B}_{\mathbf{L}, i}^{\mathcal{O}_b}(1^\eta) = 1 \cap \neg \text{bad}(\mathcal{B}_{\mathbf{L}, i})] \\ &\quad - \frac{1}{2} \mathbb{P} [\neg \text{bad}(\mathcal{B}_{\mathbf{L}, i}) \cap \neg E(\mathbf{L}, i)] \end{aligned}$$

So that

$$\begin{aligned} \text{Adv}_{\text{API}^m}^{\text{KEYSEC}}(\mathcal{A}) &= \sum_{\mathbf{L} \in \text{Act}^m} \sum_{i=1}^n \mathbb{P} [\mathcal{B}_{\mathbf{L}, i}^{\mathcal{O}_1}(1^\eta) = 1 \cap \neg \text{bad}(\mathcal{B}_{\mathbf{L}, i})] - \\ &\quad \mathbb{P} [\mathcal{B}_{\mathbf{L}, i}^{\mathcal{O}_0}(1^\eta) = 1 \cap \neg \text{bad}(\mathcal{B}_{\mathbf{L}, i})] \end{aligned}$$

which proves

$$\text{Adv}_{\text{API}^m}^{\text{KEYSEC}}(\mathcal{A}) = \sum_{\mathbf{L} \in \text{Act}^m} \sum_{i=1}^n \text{Adv}_{\text{API}^m}^{\text{KEYSEC}}(\mathcal{B}_{\mathbf{L}, i})$$

Now consider the following adversary $(\mathcal{A}_1, \mathcal{A}_2)$ in the restricted key secrecy experiment corresponding to (\mathbf{L}, i, b) : \mathcal{A}_1 simulates $\mathcal{B}_{\mathbf{L}, i}$, first sending k_i to $\mathcal{B}_{\mathbf{L}, i}$ when it makes its test query at the start of the execution, then forwarding the remaining oracle queries from $\mathcal{B}_{\mathbf{L}, i}$ to its API oracle (which is possible since $\mathcal{B}_{\mathbf{L}, i}$ follows \mathbf{L}) and sending the responses to $\mathcal{B}_{\mathbf{L}, i}$. If $\neg \text{bad}(\mathcal{B}_{\mathbf{L}, i})$ occurs then \mathcal{A}_1 does not compromise h_i and \mathcal{A}_2 receives the inputs and outputs of the oracle queries. In this case, \mathcal{A}_2 uses the queries to simulate the execution of $\mathcal{B}_{\mathbf{L}, i}$ again and outputs the same bit that $\mathcal{B}_{\mathbf{L}, i}$ outputs. Otherwise, \mathcal{A}_2 outputs a random bit. Note that sending k_i to $\mathcal{B}_{\mathbf{L}, i}$ regardless of b , but having h_i possibly point to a different key in the API, will not change the distinguishing advantage of $\mathcal{B}_{\mathbf{L}, i}$; all keys are drawn fresh at the start of each game and so the view of $\mathcal{B}_{\mathbf{L}, i}$ is the same whether the key it sees depends on b or the internal API key depends on b . Therefore it is clear that

$$\text{Adv}_{\text{API}^m, \mathbf{L}, i}^{\text{KEYSEC}}(\mathcal{A}_1, \mathcal{A}_2) = \text{Adv}_{\text{API}^m}^{\text{KEYSEC}}(\mathcal{B}_{\mathbf{L}, i})$$

So if, for each $\mathbf{L} \in \text{Act}^m$, each $1 \leq i \leq n$ and all polynomial-time adversaries $\mathcal{A}_1, \mathcal{A}_2$, the advantage $\text{Adv}_{\text{API}^m, \mathbf{L}, i}^{\text{KEYSEC}}(\mathcal{A}_1, \mathcal{A}_2)$ is negligible in η , the advantage $\text{Adv}_{\text{API}^m}^{\text{KEYSEC}}(\mathcal{A})$ must also be negligible in η for any PPT \mathcal{A} .

By a similar argument, if for each $\mathbf{L} \in \text{Act}^m$, each $1 \leq i \leq n$ and all polynomial-time adversaries $\mathcal{A}_1, \mathcal{A}_2$, the advantage $\text{Adv}_{\text{API}^m, \mathbf{L}, i}^{\text{HDLCON}}(\mathcal{A}_1, \mathcal{A}_2)$ is negligible in η , then the advantage $\text{Adv}_{\text{API}^m}^{\text{HDLCON}}(\mathcal{A})$ must also be negligible in η for any PPT \mathcal{A} . This proves the Lemma. \square

Proposition 18. *Theorem 1 follows from Theorem 2.*

Proof. The advantages $\text{Adv}_{\text{API}^m, \mathbf{L}, i}^{\text{KEYSEC}}(\mathcal{A}_1, \mathcal{A}_2)$ and $\text{Adv}_{\text{API}^m, \mathbf{L}, i}^{\text{HDLCON}}(\mathcal{A}_1, \mathcal{A}_2)$ correspond exactly to the distinguishing advantages against the key secrecy and handle consistency formulae in Theorem 2. Therefore the Proposition follows from the computational soundness theorem in [BC14] and the reduction in Lemma 17. \square

A. Basic Properties

We give here a soundness proof of the basic axioms found in Figure 4 that are not proved sound in [BC14].

Lemma 19. *The axioms in Figure 4 are sound in all computational models.*

Proof. Axioms 1 to 8 were proved sound in [BC14]. It remains to prove axioms 9 and 10. For axiom 9, we note that if $\llbracket x \rrbracket_{\eta, \rho}^\sigma = \llbracket y \rrbracket_{\eta, \rho}^\sigma$, then $\llbracket [\text{EQ}(x, y)]u(x) : w \rrbracket_{\eta, \rho}^\sigma = \llbracket [u(x)] \rrbracket_{\eta, \rho}^\sigma = \llbracket [u(y)] \rrbracket_{\eta, \rho}^\sigma$; otherwise, $\llbracket [\text{EQ}(x, y)]u(x) : w \rrbracket_{\eta, \rho}^\sigma = \llbracket [w] \rrbracket_{\eta, \rho}^\sigma$. Therefore, in either case,

$$\llbracket \vec{v}, [\text{EQ}(x, y)]u(x) : w \rrbracket_{\eta, \rho}^\sigma = \llbracket \vec{v}, [\text{EQ}(x, y)]u(y) : w \rrbracket_{\eta, \rho}^\sigma$$

and hence, if ψ is axiom 9, then for any adversary \mathcal{A} , $\text{Adv}_{\psi}^{\sim}(\mathcal{A}) = 0$.

For axiom 10, we note that $\vec{v}, x \sim \vec{v}, c$ implies $x \sim c$ by axiom 7. For the converse, we remark that, as $\llbracket [c] \rrbracket_{\eta, \rho}^\sigma$ does not depend on ρ , the computational interpretation of the formula $x \sim c$ is equivalent to the statement that $\mathbb{P}(\llbracket [x] \rrbracket_{\eta, \rho}^\sigma \neq \llbracket [c] \rrbracket_{\eta, \rho}^\sigma) = \text{negl}(\eta)$ for some negligible function negl . Hence, for all adversaries \mathcal{A} , we have

$$\mathbb{P}\left[\mathcal{A}\left(\llbracket \vec{v}, x \rrbracket_{\eta, \rho}^\sigma\right) \neq \mathcal{A}\left(\llbracket \vec{v}, c \rrbracket_{\eta, \rho}^\sigma\right)\right] \leq \text{negl}(\eta)$$

whence, for $\psi' \equiv \vec{v}, x \sim \vec{v}, c$, we have $\text{Adv}_{\psi'}^{\sim}(\mathcal{A}) \leq \text{negl}(\eta)$. \square

We now prove the lemma showing that if $\text{EQ}(x, y) \sim \text{true}$ then x and y are indistinguishable in any context.

Proof of Lemma 9. Let \vec{v} be any sequence of terms and suppose $\text{EQ}(x, y) \sim \text{true}$. We get $\vec{v}, x \sim \vec{v}, [\text{true}]x$ by axiom 1. By axiom 10, this is equivalent to $\vec{v}, x \sim \vec{v}, [\text{EQ}(x, y)]x$ and the right hand side is indistinguishable from $\vec{v}, [\text{EQ}(x, y)]y$ by axiom 9. Again axiom 10 yields $\vec{v}, x \sim \vec{v}, [\text{true}]y$ which becomes $\vec{v}, x \sim \vec{v}, y$ by axiom 1.

On the other hand, if $\vec{v}, x \sim \vec{v}, y$ for any sequence of terms \vec{v} then, in particular, $x, x \sim x, y$ and therefore $\text{EQ}(x, x) \sim \text{EQ}(x, y)$ by axiom 3. By axiom 8 and the reflexivity (axiom 4) and transitivity (axiom 6) of \sim , we find $\text{EQ}(x, y) \sim \text{true}$. \square

B. Axioms of Boolean terms

In this Section, we prove the soundness of the axioms for manipulating boolean terms.

Lemma 20. *For any $b, b_1, \dots, b_n, b'_1, \dots, b'_n$ of sort `bool`, the axioms in Figure 5 are sound in all computational models.*

Proof. The soundness of these axioms is immediate from the computational interpretations of `if _ then _ else _`, `¬`, `∧`, `∨` and `→`, and reflect the fact that the function symbols we use to operate on terms of sort `bool` perfectly capture the operators in propositional logic. \square

Lemma 21. *The case disjunction axiom 17 is sound in all computational models.*

Proof. Let $(b_1^i)_{i \in I}$ and $(b_2^i)_{i \in I}$ be partitions and assume $\vec{v}_1, b_1^i, [b_1^i]u_1 \sim \vec{v}_2, b_2^i, [b_2^i]u_2$ for all $i \in I$. Let \mathcal{A} be an adversary against $\psi \equiv \vec{v}_1, u_1 \sim \vec{v}_2, u_2$.

We consider, for each $i \in I$, the adversary \mathcal{A}_i against $\psi_i \equiv \vec{v}_1, b_1^i, [b_1^i]u_1 \sim \vec{v}_2, b_2^i, [b_2^i]u_2$ built as follows. On input \vec{x}, x_b, x_u , if $x_b = 1$, then \mathcal{A}_i returns $\mathcal{A}(\vec{x}, x_u)$. Otherwise, \mathcal{A}_i draws a random bit and returns it.

Since the families $(b_1^i)_{i \in I}$ and $(b_2^i)_{i \in I}$ perfectly partition the space of possibilities except with negligible probability, we remark that there exists a negligible function negl such that

$$\left| \text{Adv}_{\psi}^{\sim}(\mathcal{A}) - \sum_{i \in I} \text{Adv}_{\psi_i}^{\sim}(\mathcal{A}_i) \right| \leq \text{negl}(\eta)$$

and, since each $\text{Adv}_{\psi_i}^{\sim}(\mathcal{A}_i)$ is assumed to be negligible in η , this completes the proof. \square

C. Cryptographic Axioms

In this Section, we prove the soundness of the cryptographic axioms, assuming that the wrapping mechanism is a correct, secure AE-AD scheme.

Lemma 22. *The randomised (resp. deterministic) version of the axioms in Figures 6 are sound in all computational models such that:*

- *The interpretations of π_1 , π_2 and π_3 are the first, second and third projections of triples of bitstrings, respectively.*
- *The triple of algorithms $\Pi = (\text{keygen}, \text{enc}, \text{dec})$ is a correct, secure randomised (resp. deterministic) AE-AD scheme in the sense of [Rog02] (resp. [RS06]) such that an encryption or decryption of ζ returns ζ .*

Proof. We first prove the secrecy axiom. Let \vec{v} be such that $\text{enckey}(k; \vec{v})$ is satisfied. Suppose the probability, over all choices of ρ , that \mathcal{A} can distinguish $\llbracket \vec{v} \rrbracket_{\eta, \rho}^\sigma$ from $\llbracket \$k(\vec{v}) \rrbracket_{\eta, \rho}^\sigma$ is a non-negligible function of η . Then we build an adversary \mathcal{B} against the AE-AD privacy game as follows. Let $\mathcal{O}_\eta^?$ be the oracle given in the randomised AE-AD privacy game, and $\mathcal{O}_\eta^{\text{enc}}, \mathcal{O}_\eta^{d, ?}$ the encryption and challenge oracle from the deterministic AE-AD game. We will prove both soundness of the deterministic version and randomised version of the axiom at the same time as the proofs are very similar.

Then, for each term t appearing in \vec{v} , we recursively define, for the randomised case (using memoisation to ensure that we do not repeat queries):

$$\tau(t, \rho) = \begin{cases} \mathcal{O}_\eta^?(\tau(m, \rho), \tau(a, \rho), \tau(r, \rho)), & \text{if } t = \text{enc}(m, a, r, k) \\ \llbracket [f] \rrbracket(\tau(t_1, \rho), \dots, \tau(t_n, \rho)), & \text{if } t = f(t_1, \dots, t_n) \\ \llbracket [t] \rrbracket_{\eta, \rho}^\sigma & \text{otherwise.} \end{cases}$$

and for the deterministic case, we fix the message to be replaced, m , and its attribute a , and recursively define:

$$\tau^d(t, \rho) = \begin{cases} \mathcal{O}_\eta^{d,?}(\tau^d(m, \rho), \tau^d(a, \rho)), & \text{if } t = \text{enc}(m, a, r, k) \\ \mathcal{O}_\eta^{\text{enc}}(\tau^d(m', \rho), \tau^d(a, \rho)), & \text{if } t = \text{enc}(m', a', r, k), m, a \neq m', a' \\ \llbracket f \rrbracket(\tau^d(t_1, \rho), \dots, \tau^d(t_n, \rho)), & \text{if } t = f(t_1, \dots, t_n) \\ \llbracket t \rrbracket_{\eta, \rho}^\sigma & \text{otherwise.} \end{cases}$$

The sequences of bitstrings $\vec{x}_v = \tau(\vec{v}, \rho)$ and $\vec{x}_v^d = \tau^d(\vec{v}, \rho)$ are defined in the natural way.

The adversary \mathcal{B} first selects ρ uniformly at random and then constructs \vec{x}_v (\vec{x}_v^d for the deterministic case), making each oracle query at most once, i.e. each response to an oracle query is recorded by \mathcal{B} and, if the term $\text{enc}(m, a, r, k)$ appears more than once in \vec{v} , \mathcal{B} retrieves the oracle response from memory.

For the randomised case, due to the constraint $\text{enckey}(k; \vec{v})$, each nonce appears in at most one encryption under k , but this encryption can appear in multiple places in \vec{v} . Querying the encryption only once from the oracle ensures that \mathcal{B} is a nonce-respecting AE-AD adversary.

For the deterministic case, we remark that no two attributes may collide, so if attributes are syntactically different then their interpretations are different. Additionally, the premise of the axiom ensures that, with overwhelming probability $1 - p_{\text{col}}(\eta)$, the query to the challenge oracle is distinct from every query to the encryption with this attribute. It follows that, with probability $1 - p_{\text{col}}(\eta)$, \mathcal{B} is a valid adversary against the privacy game.

Having constructed \vec{x}_v (resp. \vec{x}_v^d), \mathcal{B} submits \vec{x}_v (resp. \vec{x}_v^d) to \mathcal{A} and returns $\mathcal{A}(\vec{x}_v)$ (resp. $\mathcal{A}(\vec{x}_v^d)$). We claim that the distribution of \vec{x}_v (resp. \vec{x}_v^d) is equal to the distribution of $\llbracket \vec{v} \rrbracket_{\eta, \rho}^\sigma$ or $\llbracket \$_k(\vec{v}) \rrbracket_{\eta, \rho}^\sigma$ (resp. $\llbracket \$_k(\vec{v}, m, a) \rrbracket_{\eta, \rho}^\sigma$), depending on the AE-AD oracle. Suppose there is a (computational) key k' such that $\mathcal{O}_\eta^? = \mathcal{O}_{\eta, k'}^{\text{enc}}$. Since k' is sampled randomly from $\text{keygen}(1^\eta)$, the distribution of k' is equal to the distribution of $\llbracket k \rrbracket_{\eta, \rho}^\sigma$ with ρ selected uniformly at random by \mathcal{B} . Similarly, if $\mathcal{O}_\eta^? = \mathcal{O}_\eta^{\text{enc}}$, then the oracle, on input (m, a, r) , returns a random string of the same length as $\text{enc}(m, a, r, k)$; this is exactly what is given by the interpretation of $\$_k(\text{enc}(m, a, r, k), r')$ for a fresh long nonce r' . This proves that the distribution of \vec{x}_v is equal to the distribution of $\llbracket \vec{v} \rrbracket_{\eta, \rho}^\sigma$ or $\llbracket \$_k(\vec{v}) \rrbracket_{\eta, \rho}^\sigma$. Similar arguments give the result in the deterministic case, depending on whether the challenge oracle return a genuine encryption or a random string.

We conclude that in the randomised case, with $\psi \equiv \vec{v} \sim \$_k(\vec{v})$,

$$\text{Adv}_{\psi}^{\sim}(\mathcal{A}) \leq \text{Adv}_{\Pi}^{\text{PRIV}}(\mathcal{B}).$$

In the deterministic case, we have to take into account the possibility of a collision of messages, we get, with $\psi^d \equiv \vec{v} \sim \$_k(\vec{v}, m, a)$

$$\text{Adv}_{\psi^d}^{\sim}(\mathcal{A}) \leq \text{Adv}_{\Pi}^{\text{PRIV}}(\mathcal{B}) + p_{\text{col}}.$$

Next we prove the integrity axiom. As the authenticity game is the same in the deterministic and randomised cases, we have one proof for both cases. Let x be such that $\text{enckey}(k; x)$ and the integrity axiom is false, i.e. there is a polynomial-time adversary \mathcal{A} which can distinguish between

$$\left[\text{EQ}(\text{wf}(\text{uwp}(x, k)), \text{wf}(x) \wedge \bigvee_{\{m\}_{k,a}^r \in \text{st}(x)} \text{EQ}(x, \{m\}_{k,a}^r)) \right]_{\eta, \rho}^\sigma$$

and $\llbracket \text{true} \rrbracket_{\eta, \rho}^\sigma$. Rewriting the previous statement, the bits b_1 and b_2 are not equal with non-negligible probability p , where

$$b_1 = \llbracket \text{wf}(\text{uwp}(x, k)) \rrbracket_{\eta, \rho}^\sigma$$

$$b_2 = \left[\text{wf}(x) \wedge \bigvee_{\{m\}_{k,a}^r \in \text{st}(x)} \text{EQ}(x, \{m\}_{k,a}^r) \right]_{\eta, \rho}^\sigma$$

We show that, under this assumption, there is an adversary \mathcal{B} such that $\text{Adv}_{\Pi}^{\text{AUTH}}(\mathcal{B}) = p$.

First remark that, if $b_2 = 1$ then we may replace x in b_1 by $\{m\}_{k,a}^r$ which, by correctness, gives $b_1 = 1$ (if x is well formed, then so is m by the failure propagation axiom). So if $b_1 \neq b_2$ then we must have $b_1 = 1$ and $b_2 = 0$. Recall that

$$\text{uwp}(x, k) \equiv \llbracket \text{triple}(x) \rrbracket \text{uwp}(\pi_1(x), \pi_2(x), \pi_3(x), k).$$

Therefore, $b_1 = 1$ if and only if

$$\begin{aligned} \llbracket x \rrbracket_{\eta, \rho}^\sigma &= \llbracket (\pi_1(x), \pi_2(x), \pi_3(x)) \rrbracket_{\eta, \rho}^\sigma \\ &= \left(\llbracket \pi_1(x) \rrbracket_{\eta, \rho}^\sigma, \llbracket \pi_2(x) \rrbracket_{\eta, \rho}^\sigma, \llbracket \pi_3(x) \rrbracket_{\eta, \rho}^\sigma \right) \end{aligned}$$

such that $\text{dec}(\llbracket \pi_1(x) \rrbracket_{\eta, \rho}^\sigma, \llbracket \pi_2(x) \rrbracket_{\eta, \rho}^\sigma, \llbracket \pi_3(x) \rrbracket_{\eta, \rho}^\sigma, \llbracket k \rrbracket_{\eta, \rho}^\sigma)$ succeeds. Thus x will yield a forgery in the AE-AD authenticity game.

The adversary \mathcal{B} proceeds as follows. First, \mathcal{B} selects (sufficiently many bits of) ρ at random and returns the triple $(\tau(\pi_1(x), \rho), \tau(\pi_2(x), \rho), \tau(\pi_3(x), \rho))$, where τ is as above but with $\mathcal{O}_\eta^? = \mathcal{O}_{\eta, k'}^{\text{enc}}$ and $k' \leftarrow \text{keygen}(1^\eta)$ as in the AE-AD authenticity game (making each oracle query at most once, as before). By the constraint $\text{enckey}(k; x)$, \mathcal{B} is a nonce-respecting adversary. If $b_1 = 1$ and $b_2 = 0$, which occurs with probability p , then $\tau(x, \rho)$ is not the output of an oracle query, but is such that $\text{dec}(\tau(\pi_1(x), \rho), \tau(\pi_2(x), \rho), \tau(\pi_3(x), \rho), s_k) \neq \zeta$. Thus \mathcal{B} outputs a forgery with probability p .

The soundness of the correctness axioms and the failure propagation axiom is immediate. It remains to prove the nonce length consistency axiom. Since $\text{o}_k(\vec{v})$ only differs from \vec{v} in the plaintexts encrypted under k , but the plaintexts in $\text{o}_k(\vec{v})$ and \vec{v} have the same length, and since $\$_k(\vec{v})$ only depends on the length of plaintexts encrypted under k , we have that $\$_k(\text{o}_k(\vec{v})) \sim \$_k(\vec{v})$. Note that the two sets of terms are not necessarily equal, since the long nonces chosen to replace the encryptions can differ in the two sets. The random strings used to replace the encryptions will not necessarily be the same, but

will have the same length and therefore come from the same distribution. \square

D. Consequences of the Cryptographic Axioms

In this Section we derive the weak secrecy axiom, the usable secrecy and integrity axioms and the key freshness axiom from the previous cryptographic axioms.

Lemma 23. *The axioms in Figure 7 are sound in any model for the previous axioms.*

Proof. First we prove the weak secrecy axiom. For the non-deterministic case, the result follows from the random string length consistency axioms and the strong secrecy axiom, as follows:

Assume that $\text{enckey}(k; \vec{v})$. Then we also have $\text{enckey}(k; \circ_{k,a}(\vec{v}))$. So, by the strong secrecy axiom, we have $\vec{v} \sim \$k(\vec{v})$ and $\circ_{k,a}(\vec{v}) \sim \$k(\circ_{k,a}(\vec{v}))$. By the random string length consistency axiom, we obtain $\circ_{k,a}(\vec{v}) \sim \$k(\vec{v})$. By transitivity and reflexivity of \sim , we have $\vec{v} \sim \circ_{k,a}(\vec{v})$.

For the deterministic case, we can use the same reasoning as above once we have checked that the additional premise of the strong secrecy axiom holds for \vec{v} and $\circ_{k,a}(\vec{v})$, i.e. that syntactically different plaintexts under k do not collide. As the attribute a is well-used for k in \vec{v} (and hence in $\circ_{k,a}(\vec{v})$ too), all encryptions under k with associated data a in these terms are encryptions of keys. Therefore it remains to show that syntactically different keys do not collide. This follows from the integrity axiom. Firstly, consider $x' \equiv \{x\}_{k',z}^y$. Clearly we have $\text{enckey}(k; x')$ since k does not appear in x' . Applying the integrity axiom, we have $\text{wf}(\text{uwp}(\{x\}_{k',z}^y, k)) = \text{false}$ since the disjunction on the right hand side is empty. By the definition of wf and the axioms for boolean terms, we obtain $\text{uwp}(\{x\}_{k',z}^y, k) = \text{fail}$. On the other hand, the axioms for equality show that $\text{EQ}(k, k') \rightarrow (\text{uwp}(\{x\}_{k',z}^y, k) = x)$ from which we get $\text{EQ}(k, k') \rightarrow (\text{fail} = x)$. With, for example, $x = \text{true}$ we can obtain $\text{EQ}(k, k') = \text{false}$, as required. (We actually need another axiom: $\text{EQ}(\text{true}, \text{false}) = \text{EQ}(\text{true}, \text{fail}) = \text{EQ}(\text{false}, \text{fail}) = \text{false}$, which is obviously sound in any computational model.)

The proof of the usable secrecy axiom follows intuitively the hybrid argument for proving computational soundness of static equivalence in [AR07]. Indeed, we start by replacing the encryptions under keys k such that $\text{enckey}(k; \vec{v})$ by encryptions of fresh keys, and continue by induction. In the CCSA model, this hybrid argument is simply reflected by using transitivity of \sim .

More precisely, we reason by induction on the length of the derivation of usability of k . If the length of the derivation is 1, the result holds trivially as $\text{enckey}(k, \vec{v})$ holds.

Assume that the axiom holds for any k, \vec{x} such that k can be deemed usable in \vec{x} in at most n steps. Let \vec{v} be such that k can be shown usable in \vec{v} in at most $n + 1$ steps. Let K be the set of keys k' occurring at the last step of the usability

proof of k and let $A_{k'}$ be the corresponding set of well-used attributes.

For each $k' \in K$ we have $\text{enckey}(k'; \vec{v})$ and hence $\text{enckey}(k'; \circ_{k,a}(\vec{v}))$. For each $a' \in A_{k'}$ we have $\text{wellused}(a', \vec{v}, k')$ and hence $\text{wellused}(a', \circ_{k,a}(\vec{v}), k')$. We write $A = (A_{k'})_{k' \in K}$. We write $\circ_{K,A}(\vec{v})$ as a shorthand for $\circ_{k_1, A_{k_1}} \circ \dots \circ \circ_{k_n, A_{k_n}}(\vec{v})$ with $K = \{k_1, \dots, k_n\}$. Multiple applications of the weak secrecy axiom yield

$$\vec{v} \sim \circ_{K,A}(\vec{v}) \quad \text{and} \quad \circ_{k,a}(\vec{v}) \sim \circ_{K,A}(\circ_{k,a}(\vec{v}))$$

As k can be deemed usable in at most n steps in $\circ_{K,A}(\vec{v})$, the induction hypothesis yields

$$\circ_{K,A}(\vec{v}) \sim \circ_{k,a}(\circ_{K,A}(\vec{v}))$$

We conclude the proof by transitivity of \sim and remarking that $\circ_{k,a}(\circ_{K,A}(\vec{v})) \equiv \circ_{K,A}(\circ_{k,a}(\vec{v}))$.

The usable secrecy axiom and the (first) integrity axiom together imply the usable integrity axiom.

We say that (k', a) protects k in x if there is a wrap of the form $\{k\}_{k',a}$ in $\text{st}(x)$. We define the protecting set K of k in x as the least fixed point of the protects relation containing k .

Now, k is usable in x if and only if:

- Every $(k', a) \in K$ is usable in x with a well used for k' .
- k only appears protected by keys in K or in an encryption key position in x (with nonces used correctly and proper attributes)

For a term t we denote by \tilde{t} the result of replacing any wrap $\{k'\}_{k',a}$ in $\text{st}(t)$, where $k'' \in K$, by $\{r_{k',k''}\}_{k',a}$.

Now let

$$t \equiv \text{EQ} \left(\text{wf}(\text{uwp}(x, k)), \text{wf}(x) \wedge \left(\bigvee_{\{m\}_{k,a}^r \in \text{st}(x)} \text{EQ}(x, \{m\}_{k,a}^r) \right) \right)$$

We must prove that $t \sim \text{true}$ if k is usable in x .

Since K is also the protecting set of k in t , and k is usable in t , every key in K is usable in t . Therefore by repeated applications of the usable secrecy axiom we get $t \sim \tilde{t}$.

Now, observe that we have $\text{enckey}(k; \tilde{t})$ by the definition of \tilde{t} and K . In particular, all wraps under k in $\text{st}(x)$ have a term r of sort nonce in the position of the nonce (such that $\tilde{r} \equiv r$). These two remarks show that, if $w \equiv \{m\}_{k,a}^r \in \text{st}(x)$, then $\tilde{w} \equiv \{\tilde{m}\}_{k,\tilde{a}}^r \in \text{st}(\tilde{x})$. Thus,

$$\tilde{t} \equiv \text{EQ} \left(\text{wf}(\text{uwp}(\tilde{x}, k)), \text{wf}(\tilde{x}) \wedge \left(\bigvee_{\{m\}_{k,a}^r \in \text{st}(x)} \text{EQ}(\tilde{x}, \{\tilde{m}\}_{k,\tilde{a}}^r) \right) \right)$$

and we would like to apply the integrity axiom to conclude, but *a priori* the large disjunction operator has the wrong form. Nevertheless, each $\{m\}_{k,a}^r \in \text{st}(x)$ corresponds to the term

$\{\tilde{m}\}_{k,\tilde{a}}^r \in \text{st}(\tilde{x})$. So we can index the disjuncts by this term instead. This shows that

$$\begin{aligned} \bigvee_{\{m\}_{k,a}^r \in \text{st}(x)} \text{EQ}(\tilde{x}, \{\tilde{m}\}_{k,\tilde{a}}^r) &\equiv \bigvee_{\{\tilde{m}\}_{k,\tilde{a}}^r \in \text{st}(\tilde{x})} \text{EQ}(\tilde{x}, \{\tilde{m}\}_{k,\tilde{a}}^r) \\ &\equiv \bigvee_{\{m\}_{k,a}^r \in \text{st}(\tilde{x})} \text{EQ}(\tilde{x}, \{m\}_{k,a}^r) \end{aligned}$$

Since $\text{enckey}(k; \tilde{x})$, we may use the integrity axiom to obtain $\tilde{t} \sim \text{true}$. Then $t \sim \text{true}$ by the transitivity of \sim .

Finally, the key freshness axiom follows easily from the weak secrecy axiom, and the basic axioms. \square

E. Policy Axioms

Finally, we use the properties of valid policies to derive the axioms describing key compromise.

Lemma 24. *Axioms (34) and (35) are logical consequences of the axioms in Figures 4, 5 and the policy axioms in Figure 8.*

Proof. Axiom (34). Recall that we overload \mathbf{comp} : for a handle h we take $\mathbf{comp}_q(h)$ as a shorthand for $\mathbf{comp}_q(\mathbf{a}_q(h))$.

First, for $x \in \mathcal{T} \setminus \mathcal{H}$ we have:

$$\mathbf{comp}_q(x) \equiv \bigvee_{a \in \mathcal{D}} \bigvee_{b \in \mathbf{sacr}(a)} \mathbf{cor}_q(a) \wedge \text{EQ}(x, b)$$

It is enough to remark that $\mathbf{cor}_q(a) \rightarrow \mathbf{cor}_{q'}(a) = \text{true}$ to conclude.

Now let $h \in \mathcal{H}$. If h is unused in state q , which is a property of the state, we have $\mathbf{comp}_q(h) = \text{false}$ which yields our result. Now assume that h is not unused so that $\mathbf{comp}_q(h) \equiv \mathbf{comp}_q(\mathbf{a}_q(h))$. We remark that $\mathbf{a}_{q'}(h) = \mathbf{a}_q(h)$ since the value of used handles are never modified. The result follows from the case $x \in \mathcal{T} \setminus \mathcal{H}$:

$$\mathbf{comp}_q(\mathbf{a}_q(h)) \rightarrow \mathbf{comp}_{q'}(\mathbf{a}_q(h))$$

Axiom (35). First remark that

$$\begin{aligned} &\left(\bigvee_{a \in \mathcal{D}} \bigvee_{b \in \mathbf{sacr}(a)} \mathbf{cor}_q(a) \wedge \text{EQ}(y, b) \right) \wedge \text{P}(x, y) \\ &= \left(\bigvee_{a \in \mathcal{D}} \bigvee_{b \in \mathbf{sacr}(a)} \mathbf{cor}_q(a) \wedge \text{EQ}(y, b) \wedge \text{P}(x, y) \right) \\ &= \left(\bigvee_{a \in \mathcal{D}} \bigvee_{b \in \mathbf{sacr}(a)} \mathbf{cor}_q(a) \wedge \text{EQ}(y, b) \wedge \text{P}(x, b) \right) \end{aligned}$$

By validity of the policy and transitivity of the sacrifice function, we get for any a

$$\begin{aligned} &\bigvee_{b \in \mathbf{sacr}(a)} \text{EQ}(y, b) \wedge \text{P}(x, b) \\ &= \bigvee_{b \in \mathbf{sacr}(a)} \text{EQ}(y, b) \wedge \text{P}(x, b) \wedge \bigvee_{c \in \mathbf{sacr}(b)} \text{EQ}(x, c) \\ &= \bigvee_{b, c \in \mathbf{sacr}(a)} \text{EQ}(y, b) \wedge \text{P}(x, b) \wedge \text{EQ}(x, c) \end{aligned}$$

Using the definition of \mathbf{comp} and the two previous equalities, we get

$$\mathbf{comp}_q(y) \wedge \text{P}(x, y) = \mathbf{comp}_q(y) \wedge \text{P}(x, y) \wedge \mathbf{comp}_q(x)$$

The lemma then follows from Axiom 16. \square

APPENDIX D PROOF OF THE MAIN THEOREM

Before tackling the main Theorem, we prove that the frames ϕ^* and ϕ^\dagger , which are *a priori* very complicated sequences of terms, can be rewritten in a relatively simple form. First, we must introduce some notation:

In an API execution with initial state q_0 , if $\mathbf{k}_{q_0}(h) = k$ then we will write a_k for $\mathbf{a}_{q_0}(h)$ and often call this *the attribute* of k . It is well-defined given the assumption from Section II that all initial configurations are honest. By convention, $a_{\text{fail}} \equiv \text{fail}$. Finally we write

$$K_0 = \{k \in \mathcal{K} \mid \exists h \in \mathcal{H}, \mathbf{k}_{q_0}(h) = k\}.$$

The following definition is used to identify sets of terms in which keys in K_0 appear only as expected: if they are in a plaintext position then they are protected by a properly-constructed wrap or known to have a compromised attribute; if they appear in an unwrap position (which cannot be simplified using the integrity and correctness axioms) then they are known to have a compromised attribute; and if they appear in a wrapping position then the key's attribute is internal and, either the wrap is of a key from K_0 and its corresponding attribute, or the associated data of the wrap is a compromised attribute. The only caveat is that, if (in the $*$ case), the key pointed to by h^* is initially given to the adversary, then there is an element $k^* \in K_0$ which appears in plaintext without initially having a compromised attribute. So the condition on keys appearing in plaintext positions can only be assured for *internal* keys in this case (since k^* is an arbitrary external key). Therefore we have to have two slightly different versions of this definition depending on whether h^* points to k^* or k^\dagger :

Definition 25. Let X be a set of terms. We say X is $(q_0, q, F)^\dagger$ -normal if, for each $k \in K_0$, the following holds:

- 1) If k appears in a plaintext position $p.\alpha$ in X , then
 - a) either $X|_p = \{k\}_{k', a_k}^r$ where $k' \in K_0$, $\text{P}(a_k, a_{k'}) = \text{true}$ and r has sort nonce
 - b) or $(F \rightarrow \mathbf{comp}_q(a_k)) = \text{true}$
- 2) If k appears in an unwrap key position in X , then $(F \rightarrow \mathbf{comp}_q(a_k)) = \text{true}$

- 3) If $\{u\}_{k,v}^w \in \text{st}(X)$ then $a_k \notin \mathcal{E}$, v has sort attribute and
- either $u \equiv k', v \equiv a_{k'}$ for some $k' \in K_0$
 - or $(F \rightarrow \mathbf{comp}_q(v)) = \text{true}$

If Point 1) holds for all $k \in K_0$ with $a_k \notin \mathcal{E}$ (but the other Points hold for all $k \in K_0$), then we say X is $(q_0, q, F)^*$ -normal).

Lemma 26. *Say P is valid with respect to sacr . If X is $(q_0, q, F)^\dagger$ -normal and $k \in K_0$ has a_k uncompromised in q , then k is usable in X . If X is $(q_0, q, F)^*$ -normal and $k \in K_0$ has a_k uncompromised in q and $a_k \notin \mathcal{E}$, then k is usable in X .*

Proof. The Lemma follows immediately from the definition of usable keys. \square

The next definition identifies sequences of boolean terms that partition the space of possibilities in a neat way (in the sense of normality as defined above) and give sufficient information to determine whether or not any given handle or attribute is compromised.

Definition 27. Let $q \in \mathbf{Q}$ and I_q an indexing set. We say the sequence of boolean terms $(\psi_q^i)_{i \in I_q}$ is a $(q_0, q)^\dagger$ -valid partition (respectively, a $(q_0, q)^*$ -valid partition) if the following hold:

- $\bigvee_{i \in I_q} \psi_q^i = \text{true}$
- For each $i, j \in I_q$ with $i \neq j$, we have $\psi_q^i \wedge \psi_q^j \sim \text{false}$.
- For each $i \in I_q$, ψ_q^i is $(q_0, q, \psi_q^i)^\dagger$ -normal (resp., $(q_0, q, \psi_q^i)^*$ -normal)
- For each $i \in I_q$ and $x \in \mathcal{H} \cup \mathcal{D}$, either $(\psi_q^i \rightarrow \mathbf{comp}_q(x)) = \text{true}$ or $(\psi_q^i \rightarrow \neg \mathbf{comp}_q(x)) = \text{true}$.

Remark 28. Due to Axiom 34, if there is a transition from state q to state q' then a $(q_0, q, F)^\dagger$ -normal set X is also $(q_0, q', F)^\dagger$ -normal set and a $(q_0, q, F)^*$ -normal set X is also $(q_0, q', F)^*$ -normal set. However a $(q_0, q)^\dagger$ -valid partition is not necessarily a $(q_0, q')^\dagger$ -valid partition (and similarly for $*$) since $\psi_q^i \rightarrow \neg \mathbf{comp}_q(x)$ does not imply $\psi_{q'}^i \rightarrow \neg \mathbf{comp}_{q'}(x)$.

Now we formalise rewriting the frame in a simple form, giving six invariants that are true of the terms output by any API execution. The main idea is that the usable integrity axiom and Lemma 26 will together ensure that, if the adversary's free function $g(\phi)$ is successfully unwrapped using an uncompromised key k , then $g(\phi)$ must be semantically equal to a previously created wrap under k . This partitions the possibility space for $g(\phi)$ into the different wraps under k appearing in ϕ and, within each section of this partition, i.e. each choice of $g(\phi)$, the term containing the unwrap can be simplified by the correctness axiom. So one can think of the booleans F^i mentioned in the invariants as determining which of the possible wraps was unwrapped at each stage.

Invariant 1 says that we can choose the F^i to have a normal form that also enables one to deduce whether or not a given handle or attribute is compromised.

Invariant 2 says the frame can be rewritten in a normal form under the conditions of the F^i (by rewriting unwraps).

Invariant 3 says that, under the conditions of the F^i , if a handle is uncompromised then it points to a key from K_0 and the attribute of that key, or the pair (fail, fail). That is, handles not known to be compromised cannot point to keys imported by the adversary.

Invariant 4 says that, under the conditions of the F^i , if a handle is compromised then the handle points to a term in normal form whose second projection is a compromised attribute.

Invariant 5 says that each wrap appearing in the rewritten terms uses valid randomness: names of sort nonce not used anywhere other than in the nonce position of this wrap.

Finally, invariant 6 says that the conditions F^i in the \dagger case can be obtained by the F^i in the $*$ case (and similarly for the corresponding rewritten frame, etc.), just by replacing every occurrence of k^\dagger by k^* .

Proposition 29. *Under the hypotheses of Theorem 2, for the executions*

$$(q_0^*, (k^*, \phi_0)) \rightarrow \dots \rightarrow (q^*, \phi^*)$$

$$(q_0^\dagger, (k^\dagger, \phi_0)) \rightarrow \dots \rightarrow (q^\dagger, \phi^\dagger)$$

there exist an indexing set I_{q^\dagger} and a sequence of booleans $(F_{q^\dagger}^i)_{i \in I_{q^\dagger}}$ (respectively, an indexing set I_{q^*} and a sequence of booleans $(F_{q^*}^i)_{i \in I_{q^*}}$) such that:

- $(F_{q^\dagger}^i)_{i \in I_{q^\dagger}}$ is a $(q_0^\dagger, q^\dagger)^\dagger$ -valid partition (resp., $(F_{q^*}^i)_{i \in I_{q^*}}$ is a $(q_0^*, q^*)^*$ -valid partition)
- There exists a family $(\phi_{q^\dagger}^i)_{i \in I_{q^\dagger}}$ (resp., $(\phi_{q^*}^i)_{i \in I_{q^*}}$) of frames, such that for each $i \in I_{q^\dagger}$ (resp., $i \in I_{q^*}$):
 - $[F_{q^\dagger}^i] \phi^\dagger = [F_{q^\dagger}^i] \phi_{q^\dagger}^i$ (resp., $[F_{q^*}^i] \phi^* = [F_{q^*}^i] \phi_{q^*}^i$)
 - $\phi_{q^\dagger}^i$ is $(q_0^\dagger, q^\dagger, F_{q^\dagger}^i)^\dagger$ -normal (resp., $\phi_{q^*}^i$ is $(q_0^*, q^*, F_{q^*}^i)^*$ -normal)
- There exists a family $(\kappa_{q^\dagger}^i)_{i \in I_{q^\dagger}}$ (resp., $(\kappa_{q^*}^i)_{i \in I_{q^*}}$) of maps $\mathcal{H} \rightarrow K_0 \cup \{\text{fail}\}$ such that, for each $i \in I_{q^\dagger}$ (resp., $i \in I_{q^*}$) and each $h \in \mathcal{H}$ such that $F_{q^\dagger}^i \rightarrow \neg \mathbf{comp}_{q^\dagger}(h) = \text{true}$,
$$[F_{q^\dagger}^i] \mathbf{val}_{q^\dagger}(h) = [F_{q^\dagger}^i] \left(\kappa_{q^\dagger}^i(h), a_{\kappa_{q^\dagger}^i(h)} \right)$$

respectively,

$$[F_{q^*}^i] \mathbf{val}_{q^*}(h) = [F_{q^*}^i] \left(\kappa_{q^*}^i(h), a_{\kappa_{q^*}^i(h)} \right)$$
- There exist a family $(\tau_{q^\dagger}^i)_{i \in I_{q^\dagger}}$ (resp., $(\tau_{q^*}^i)_{i \in I_{q^*}}$) of maps $\mathcal{H} \rightarrow \mathcal{T} \times \mathcal{D}$ such that, for each $i \in I_{q^\dagger}$ (resp., $i \in I_{q^*}$) and each $h \in \mathcal{H}$ such that $F_{q^\dagger}^i \rightarrow \mathbf{comp}_{q^\dagger}(h) = \text{true}$

- a) $\left[F_{q^\dagger}^i\right] \mathbf{val}_{q^\dagger}(h) = \left[F_{q^\dagger}^i\right] \tau_{q^\dagger}^i(h)$
resp., $\left[F_{q^*}^i\right] \mathbf{val}_{q^*}(h) = \left[F_{q^*}^i\right] \tau_{q^*}^i(h)$
- b) if $\tau_{q^\dagger}^i(h) \equiv (t, a)$ (*resp.* $\tau_{q^*}^i(h) \equiv (t, a)$) then
- t is $\left(q_0^\dagger, q^\dagger, F_{q^\dagger}^i\right)^\dagger$ -normal (*resp.* $\left(q_0^*, q^*, F_{q^*}^i\right)^*$ -normal)
 - $\left(F_{q^\dagger}^i \rightarrow \mathbf{comp}_{q^\dagger}(a)\right) = \mathbf{true}$
resp., $\left(F_{q^*}^i \rightarrow \mathbf{comp}_{q^*}(a)\right) = \mathbf{true}$
- 5) For each $i \in I_{q^\dagger}$ (*resp.*, $i \in I_{q^*}$) and each $x \equiv \{u\}_{v,w}^r \in \text{st}\left(F_{q^\dagger}^i, \phi_{q^\dagger}^i, \tau_{q^\dagger}^i(h)\right)_{h \in \mathcal{H}}$
resp., $\{u\}_{v,w}^r \in \text{st}\left(F_{q^*}^i, \phi_{q^*}^i, \tau_{q^*}^i(h)\right)_{h \in \mathcal{H}}$
- r is a name of sort nonce
 - r appears nowhere in $\text{st}\left(F_{q^\dagger}^i, \phi_{q^\dagger}^i, \tau_{q^\dagger}^i(h)\right)_{h \in \mathcal{H}}$ other than in the position of the randomness in x .
- 6) $I_{q^\dagger} = I_{q^*}$ such that for every $i \in I_{q^\dagger}$
- $F_{q^*}^i = F_{q^\dagger}^i [k^*/k^\dagger]$
 - $\phi_{q^*}^i = \phi_{q^\dagger}^i [k^*/k^\dagger]$
 - for every $h \in \mathcal{H}$ we have $\kappa_{q^*}^i(h) = \kappa_{q^\dagger}^i(h) [k^*/k^\dagger]$
 - for every $h \in \mathcal{H}$ we have $\tau_{q^*}^i(h) = \tau_{q^\dagger}^i(h) [k^*/k^\dagger]$
- (where, for any terms t, t' and any sequence of terms \vec{v} , we write $\vec{v}[t'/t]$ for the sequence of terms obtained by replacing every occurrence of t in \vec{v} by t' .)

Proof of Proposition 29. The proof is by induction and considering the effect of each possible fully-specified action.

We first prove the invariants in the \dagger case, and then apply invariant 6 to obtain the corresponding invariants in the $*$ case.

Note that the base case is obvious. Let $\Delta \in \{*, \dagger\}$. We set $I_{q_0^\Delta} = \{-\}$, $F_{q_0^\Delta}^- = \mathbf{true}$, $\phi_{q_0^\Delta}^- = (k^*, \phi_0)$, $\kappa_{q_0^\Delta}^- = \mathbf{val}_{q_0^\Delta}$ and $\tau_{q_0^\Delta}^-(h) = (\mathbf{fail}, \mathbf{bad})$ for all $h \in \mathcal{H}$. Then invariants 1, 2, 3, 4, 5 and 6 are immediate.

Now suppose the result holds for the execution $\left(q_0^\dagger, (k^*, \phi_0)\right) \rightarrow \dots \rightarrow (q^\dagger, \phi^\dagger)$ and let the action \mathbf{A} induce the transition $(q^\dagger, \sigma^\dagger, \phi^\dagger) \rightarrow (q^\Delta, \sigma^\Delta, \phi^\Delta)$.

Action $\mathbf{A} \equiv \mathbf{W}(h_1, h_2)$.

We take $I_{q^\Delta} \equiv I_{q^\dagger}$ and $(F_{q^\Delta}^i)_{i \in I_{q^\Delta}} \equiv (F_{q^\dagger}^i)_{i \in I_{q^\dagger}}$. Since no new attributes or handles are compromised by the wrap action we have $\mathbf{comp}_{q^\Delta}(x) = \mathbf{comp}_{q^\dagger}(x)$ which means that a $(q_0, q^\dagger)^\dagger$ -valid partition is a $(q_0, q^\Delta)^\dagger$ -valid partition, proving invariant 1.

Now, we prove invariant 2. We have $\phi^\Delta \equiv \phi^\dagger; t$ where

$$t \equiv [\mathbf{P}(\mathbf{a}_{q^\dagger}(h_1), \mathbf{a}_{q^\dagger}(h_2))] \{ \mathbf{k}_{q^\dagger}(h_1) \}_{\mathbf{k}_{q^\dagger}(h_2), \mathbf{a}_{q^\dagger}(h_1)}^r.$$

By the axioms in Section V-A, we then have $\left[F_{q^\Delta}^i\right] \phi^\Delta \equiv \left[F_{q^\dagger}^i\right] \phi^\dagger; \left[F_{q^\dagger}^i\right] t$. By 1 in the induction hypothesis, each $F_{q^\dagger}^i$ tells us whether or not h_1 and h_2 are compromised in q^\dagger . We consider each case in turn. Suppose i is such that h_1 and h_2 are both uncompromised in q^\dagger . Write $k_1 \equiv \kappa_{q^\dagger}^i(h_1)$ and $k_2 \equiv \kappa_{q^\dagger}^i(h_2)$. By invariant 3 of the induction hypothesis, we

may replace $\mathbf{val}_{q^\dagger}(h_1)$ and $\mathbf{val}_{q^\dagger}(h_2)$ in t by (k_1, a_{k_1}) and (k_2, a_{k_2}) . Therefore, by the correctness of π_1 and π_2 ,

$$\left[F_{q^\Delta}^i\right] \phi^\Delta \equiv \left[F_{q^\dagger}^i\right] \phi_{q^\dagger}^i; [\mathbf{P}(a_{k_1}, a_{k_2})] \{k_1\}_{k_2, a_{k_1}}^r.$$

Either $\mathbf{P}(a_{k_1}, a_{k_2}) = \mathbf{true}$ or $\mathbf{P}(a_{k_1}, a_{k_2}) = \mathbf{false}$ and this is determined by h_1, h_2 and i . In either case, the resulting $\phi_{q^\Delta}^i$ is $\left(q_0^\dagger, q^\Delta, F_{q^\Delta}^i\right)^\dagger$ -normal since $\phi_{q^\dagger}^i$ is.

Say h_1 is uncompromised in q^\dagger but h_2 is compromised. By the definition of compromised handles and the transitivity of \mathbf{sacr} , $\left[F_{q^\Delta}^i\right] \phi^\Delta = \left[F_{q^\dagger}^i\right] \phi_{q^\dagger}^i; \mathbf{fail}$, which obviously satisfies the conditions in invariant 2.

Next suppose h_1 is compromised in q^\dagger but h_2 is uncompromised. By replacing $\mathbf{val}_{q^\dagger}(h_1)$ by $\tau_{q^\dagger}^i(h_1)$ and $\mathbf{val}_{q^\dagger}(h_2)$ by (k_2, a_{k_2}) due to invariants 3 and 4, we find that

$$\left[F_{q^\Delta}^i\right] \phi^\Delta = \left[F_{q^\dagger}^i\right] \phi_{q^\dagger}^i; [\mathbf{P}(t_2, a_{k_2})] \{t_1\}_{k_2, t_2}^r$$

where $t_1 \equiv \pi_1(\tau_{q^\dagger}^i(h_1))$ and $t_2 \equiv \pi_2(\tau_{q^\dagger}^i(h_1))$. Normality is almost immediate from the normality of $\tau_{q^\dagger}^i(h_1)$, except the requirement that a key appearing in the position of a wrapping key does not have an external attribute. There is nothing to prove if $a_{k_2} \notin \mathcal{E}$ and, if not, the validity of \mathbf{P} gives $\mathbf{P}(\pi_2(\tau_{q^\dagger}^i(h_1)), a_{k_2}) = \mathbf{false}$ and $\left[F_{q^\Delta}^i\right] \phi^\Delta = \left[F_{q^\dagger}^i\right] \phi_{q^\dagger}^i; \mathbf{fail}$, which is normal.

Finally for invariant 2, if both handles are compromised in q^\dagger then we replace $\mathbf{val}_{q^\dagger}(h_1)$ and $\mathbf{val}_{q^\dagger}(h_2)$ by $\tau_{q^\dagger}^i(h_1)$ and $\tau_{q^\dagger}^i(h_2)$, obtaining a sequence of terms that is immediately normal, thanks to the normality of $\tau_{q^\dagger}^i(h_1)$ and $\tau_{q^\dagger}^i(h_2)$.

Invariants 3 and 4 for this action are immediate from the induction hypothesis since, by construction, $\mathbf{val}_{q^\Delta} = \mathbf{val}_{q^\dagger}$, $\left(q_0, q^\dagger, F_{q^\dagger}^i\right)^\dagger$ -normality implies $\left(q_0, q^\Delta, F_{q^\Delta}^i\right)^\dagger$ -normality and handles and attributes are compromised in q^Δ if and only if they are compromised in q^\dagger . Invariant 5 is also immediate by construction since any new wrap added to the frame is produced with fresh, valid randomness.

Action $\mathbf{A} \equiv \mathbf{C}(h)$.

Again set $I_{q^\Delta} \equiv I_{q^\dagger}$ and $(F_{q^\Delta}^i)_{i \in I_{q^\Delta}} \equiv (F_{q^\dagger}^i)_{i \in I_{q^\dagger}}$. As before, $\mathbf{comp}_{q^\dagger}(x)$ implies $\mathbf{comp}_{q^\Delta}(x)$ and so $\left(q_0^\dagger, q^\dagger, F_{q^\dagger}^i\right)^\dagger$ -normality implies $\left(q_0^\dagger, q^\Delta, F_{q^\Delta}^i\right)^\dagger$ -normality. This shows that the first three conditions in the definition of a $(q_0, q^\Delta)^\dagger$ -valid partition are satisfied by $(F_{q^\Delta}^i)_{i \in I_{q^\Delta}}$. The final condition is more complicated because it is likely that new attributes (and therefore new handles) are compromised in q^Δ by \mathbf{A} .

If h is compromised in q^\dagger , then no new handles and attributes are compromised by the corruption. Therefore $(F_{q^\dagger}^i)_{i \in I_{q^\dagger}}$ does indeed give the compromised handles and attributes in q^Δ .

Suppose instead that h is uncompromised in q^\dagger so that $\left[F_{q^\dagger}^i\right] \mathbf{val}_{q^\dagger}(h) = \left[F_{q^\dagger}^i\right] (k, a_k)$ where $k = \kappa_{q^\dagger}^i(h)$. Let $x \in$

$\mathcal{H} \cup \mathcal{D}$. If x is compromised in q^\dagger then x is compromised in q^Δ . Suppose instead that x is uncompromised in q^\dagger . Say x is a handle, h' . Then we have $\left[F_{q^\dagger}^i\right] \mathbf{val}_{q^\dagger}(h') = \left[F_{q^\dagger}^i\right](k', a_{k'})$ where $k' = \kappa_{q^\dagger}^i(h')$. Then x is not compromised in q^\dagger if and only if $a_{k'} \in \mathbf{sacr}(a_k)$. Similarly, if $x \equiv a' \in \mathcal{D}$ then x is compromised in q^\dagger if and only if $a' \in \mathbf{sacr}(a_k)$. This shows that the fourth condition of a $(q_0^\dagger, q^\Delta)^\dagger$ -valid partition is satisfied by $(F_{q^\Delta}^i)_{i \in I_{q^\Delta}}$.

For invariant 2, first recall that $\phi^\Delta \equiv \phi^\dagger; \mathbf{k}_{q^\dagger}(h)$. For each i , $F_{q^\dagger}^i$ either implies that h is compromised in q^\dagger or that h is uncompromised in q^\dagger . In the first case we have $\left[F_{q^\Delta}^i\right] \phi^\Delta = \left[F_{q^\dagger}^i\right] \left(\phi_{q^\dagger}^i; \pi_1 \left(\tau_{q^\dagger}^i(h)\right)\right)$ and in the second case we have $\left[F_{q^\Delta}^i\right] \phi^\Delta = \left[F_{q^\dagger}^i\right] \left(\phi_{q^\dagger}^i; \pi_1 \left(\kappa_{q^\dagger}^i(h)\right)\right)$ where normality in both cases is immediate from the induction hypothesis.

For invariant 3, we remark that for any $h' \in \mathcal{H}$, $(\neg \mathbf{comp}_{q^\Delta}(h') \rightarrow \neg \mathbf{comp}_{q^\dagger}(h')) = \text{true}$ so we can directly apply invariant 3 of the induction hypothesis.

For invariant 4, we have to fix the value of $\tau_{q^\Delta}^i(h')$ for each $h' \in \mathcal{H}$. If h is compromised in q^\dagger then $\mathbf{val}_{q^\Delta}(h') \equiv \mathbf{val}_{q^\dagger}(h')$ by definition and we can directly apply invariant 4 of the induction hypothesis. For the latter case, invariant 3 of the induction hypothesis gives us $\left[F_{q^\dagger}^i\right] \mathbf{val}_{q^\dagger}(h') = \left[F_{q^\dagger}^i\right](k', a_{k'})$ where $k' \equiv \kappa_{q^\dagger}^i(h')$, so we take $\tau_{q^\dagger}^i(h') = (k', a_{k'})$, which is obviously normal. Note that $\pi_2 \left(\tau_{q^\dagger}^i(h')\right) = a_{k'}$ which, by the same reasoning as in the proof of invariant 1 for this action, satisfies $\left(F_{q^\dagger}^i \rightarrow a_{k'} \in \mathbf{sacr}(a_k)\right) = \text{true}$ and hence $\left(F_{q^\dagger}^i \rightarrow \mathbf{comp}_q(a_{k'})\right) = \text{true}$ as required for the final condition in invariant 4.

Invariant 5 is immediate, since we have not introduced any new wraps in any terms.

Action A $\equiv \mathbf{U}(h)$.

This is the most complicated action to prove because the $(F_{q^\Delta}^i)_{i \in I_{q^\Delta}}$ need to take care of the different possible results of unwrapping the adversary's function $g_{q^\dagger}(\phi^\dagger)$. Therefore we cannot simply reuse I_{q^\dagger} and $(F_{q^\dagger}^i)_{i \in I_{q^\dagger}}$ as we did for the previous actions. Instead we take $I_{q^\Delta} = \bigcup_{i \in I_{q^\dagger}} J_i$ where the index sets J_i will depend on the properties of $F_{q^\dagger}^i$.

Say $F_{q^\dagger}^i \rightarrow \mathbf{comp}_{q^\dagger}(h) = \text{true}$. Let $\tau_{q^\dagger}^i(h) = (t, a_h)$. Then take $J_i = \{(i, a) \mid a \in \mathcal{D}, a \in \mathbf{sacr}(a_h)\} \cup \{(i, \perp)\}$ and

$$F_{q^\Delta}^{(i,a)} \equiv F_{q^\dagger}^i \wedge \mathbf{wf}(u') \wedge \mathbf{EQ} \left(a, \pi_2(g(\phi_{q^\dagger}^i)) \right)$$

$$F_{q^\Delta}^{(i,\perp)} \equiv F_{q^\dagger}^i \wedge \left(\neg \mathbf{wf}(u') \vee \bigwedge_{a \in \mathbf{sacr}(a_h)} \left(\neg \mathbf{EQ} \left(a, \pi_2(g(\phi_{q^\dagger}^i)) \right) \right) \right)$$

where g is a shorthand for g_{n-1} , n is the number of transitions from q_0^\dagger to q^\dagger and $u' \equiv \text{uwp} \left(g \left(\phi_{q^\dagger}^i \right), t \right)$.

We check Points 2 to 5 for the (i, a) in such J_i .

By definition,

$$\phi^\Delta \equiv \phi^\dagger; [\mathbf{wf}(u) \wedge \mathbf{P}(\pi_2(g(\phi^\dagger)), \mathbf{a}_{q^\dagger}(h))] h_f$$

where $u \equiv \text{uwp} \left(g \left(\phi^\dagger \right), \mathbf{k}_{q^\dagger}(h) \right)$ and $h_f \equiv \mathbf{freshhdl}(q^\dagger)$.

By invariants 2 and 4, under the conditions of $F_{q^\dagger}^i$ we may replace any occurrence of ϕ^\dagger with $\phi_{q^\dagger}^i$ and any occurrence of $\mathbf{val}_{q^\dagger}(h)$ with $\tau_{q^\dagger}^i(h)$. Then it is easy to show that

$$\left[F_{q^\Delta}^{(i,a)}\right] \phi^\Delta = \left[F_{q^\dagger}^i\right] \left(\phi_{q^\dagger}^i; [\mathbf{wf}(u') \wedge \mathbf{P}(a, a_h)] (h_f) \right)$$

where by convention $\mathbf{P}(\perp, _) \equiv \text{false}$. The only non-trivial case of proving the above is when $a = \perp$, but it follows from the policy axioms. Now, since compromise in q^\dagger implies compromise in q^Δ , we have that $(q_0^\dagger, q^\dagger, F_{q^\dagger}^i)^\dagger$ -normality implies $(q_0^\dagger, q^\Delta, F_{q^\dagger}^i)^\dagger$ -normality, so we may take

$$\phi_{q^\Delta}^{(i,a)} \equiv \phi_{q^\dagger}^i; [\mathbf{wf}(u') \wedge \mathbf{P}(a, a_h)] (h_f)$$

which satisfies the conditions of invariant 2 thanks to the normality of $\phi_{q^\dagger}^i$ and $\tau_{q^\dagger}^i(h)$.

Invariants 3 and 4 are immediate for all handles $h' \neq h_f$ since $\mathbf{val}_{q^\Delta}(h') \equiv \mathbf{val}_{q^\dagger}(h')$ in this case and $\mathbf{comp}_{q^\Delta}(h')$ is logically equivalent to $\mathbf{comp}_{q^\dagger}(h')$, so we may take $\kappa_{q^\Delta}^{(i,a)}(h') \equiv \kappa_{q^\dagger}^i(h')$ and $\tau_{q^\Delta}^{(i,a)}(h') \equiv \tau_{q^\dagger}^i(h')$.

We prove invariant 3 for h_f . Put $\kappa_{q^\Delta}^{(i,a)}(h_f) \equiv \text{fail}$. Intuitively, this is because h_f is considered compromised in q^Δ whether or not the unwrap succeeds, since h is compromised in q^\dagger . So the conditions in invariant 3 will evaluate to false. Formally, for any term x , we have

$$\left[\neg \mathbf{comp}_{q^\Delta}(h_f) \wedge F_{q^\Delta}^{(i,a)} \right] x = \text{fail}.$$

Since we may take $x \equiv \mathbf{val}_{q^\Delta}(h_f)$ or $x \equiv (\text{fail}, \text{fail})$, the condition of invariant 3 is satisfied.

Next, we prove invariant 4 for h_f . By definition,

$$\left[F_{q^\Delta}^{(i,a)}\right] \mathbf{val}_{q^\Delta}(h_f) = \left[F_{q^\Delta}^{(i,a)}\right] \left([\mathbf{wf}(u) \wedge \mathbf{P}(t_2, \mathbf{a}_{q^\dagger}(h))] (u, t_2) : (\text{fail}, \text{bad}) \right)$$

where $t_2 \equiv \pi_2 \left(g \left(\phi_{q^\dagger}^i \right) \right)$. As before we may replace u by u' and $\mathbf{a}_{q^\dagger}(h)$ by a_h , so for $a \neq \perp$ the above is equal to

$$\left[F_{q^\Delta}^{(i,a)}\right] \left([\mathbf{P}(a, a_h)] (u', a) : (\text{fail}, \text{bad}) \right).$$

Therefore, for $a \in \mathbf{sacr}(a_h)$ such that $\mathbf{P}(a, a_h) = \text{true}$ we put $\tau_{q^\Delta}^{(i,a)}(h_f) = (u', a)$ and for $a \in \mathbf{sacr}(a_h)$ such that $\mathbf{P}(a, a_h) = \text{false}$ we put $\tau_{q^\Delta}^{(i,a)}(h_f) = (\text{fail}, \text{bad})$. For $a = \perp$, either the unwrap is not well-formed or the policy will reject the unwrap, so we may set $\tau_{q^\Delta}^{(i,\perp)}(h_f) = (\text{fail}, \text{bad})$. Since all keys from K_0 appearing in $\tau_{q^\Delta}^{(i,a)}(h_f)$ must appear in $\phi_{q^\dagger}^i$ or $\tau_{q^\dagger}^i(h)$, it follows from the induction hypothesis that $\tau_{q^\Delta}^{(i,a)}(h_f)$ is $(q_0^\dagger, q^\dagger, F_{q^\dagger}^i)^\dagger$ -normal and hence $(q_0^\dagger, q^\Delta, F_{q^\Delta}^{(i,a)})^\dagger$ -normal,

as required for Part 4(b)i. For Part 4(b)ii, recall that $\mathbf{comp}_{q^\dagger}(h) = \mathbf{comp}_{q^\dagger}(h) \wedge \mathbf{comp}_{q^\Delta}(h_f)$ so that $\left(F_{q^\dagger}^i \rightarrow \mathbf{comp}_{q^\dagger}(h) \wedge \mathbf{comp}_{q^\Delta}(h_f)\right) = \text{true}$ and hence $\left(F_{q^\Delta}^{(i,a)} \rightarrow \mathbf{comp}_{q^\Delta}(h_f)\right) = \text{true}$ as required.

Invariant 5 is immediate since there are no wraps in $\text{st}(F_{q^\Delta}^i, \phi_{q^\Delta}^i, \tau_{q^\Delta}^i)_{h \in \mathcal{H}}$ that did not appear in $\text{st}(F_{q^\dagger}^i, \phi_{q^\dagger}^i, \tau_{q^\dagger}^i)_{h \in \mathcal{H}}$.

Now we suppose $F_{q^\dagger}^i \rightarrow \neg \mathbf{comp}_{q^\dagger}(h) = \text{true}$. Set $\kappa_{q^\dagger}^i(h) \equiv k$ so that $\left[F_{q^\dagger}^i\right] \mathbf{val}_{q^\dagger}(h) = \left[F_{q^\dagger}^i\right](k, a_k)$. Then let W be the set of wraps appearing in $\text{st}(\phi_{q^\dagger}^i)$ such that k is in the wrapping key position. By invariant 5, each element $w \in W$ uses a different term for its randomness and, by the normality of $\phi_{q^\dagger}^i$, if $a_k \in \mathcal{E}$ then W is empty.

Then define $J_i \equiv \{(i, 1), (i, 2, w) \mid w \in W\}$. Indices of the form $(i, 1)$ will be used to identify when the unwrap fails. Indices of the form $(i, 2, w)$ will be used to identify when the adversary is trying to unwrap $w \in W$. We will use the usable integrity axiom to prove that these are the only two possibilities. We define

$$F_{q^\dagger}^{(i,1)} \equiv \neg \text{wf}(g(\phi_{q^\Delta}^i)) \vee \bigwedge_{w \in W} \neg \text{EQ}(w, g(\phi_{q^\Delta}^i))$$

$$F_{q^\dagger}^{(i,2,w)} \equiv \text{wf}(g(\phi_{q^\Delta}^i)) \wedge \text{EQ}(w, g(\phi_{q^\Delta}^i))$$

By invariant 5 and normality of $\phi_{q^\Delta}^i$, k is usable in $\phi_{q^\Delta}^i$. By the usable integrity axiom, we get

$$F_{q^\dagger}^{(i,1)} = \neg \text{wf}(uwp(g(\phi_{q^\Delta}^i), k))$$

In the case $j = (i, 1)$, we set $\kappa_{q^\dagger}^j(\mathbf{freshhdl}(q)) = \text{fail}$, for every $h \not\equiv \mathbf{freshhdl}(q)$ $\kappa_{q^\dagger}^j(h) = \kappa_{q^\Delta}^i(h)$ $\phi_{q^\dagger}^j = \phi_{q^\Delta}^i$, and $\tau_{q^\dagger}^j = \tau_{q^\Delta}^i$ invariants 2, 3, 4, 5 follow.

In the case $j = (i, 2, w)$, by $(q_0^\dagger, q^\Delta, F_{q^\Delta})^\dagger$ normality of $\phi_{q^\Delta}^i$ yields

- Either $w = \{k'\}_{k, a_{k'}}$, we set $\kappa_{q^\dagger}^j(\mathbf{freshhdl}(q)) = k'$. For every $h \not\equiv \mathbf{freshhdl}(q)$, put $\kappa_{q^\dagger}^j(h) = \kappa_{q^\Delta}^i(h)$. Also set $\tau_{q^\dagger}^j = \tau_{q^\Delta}^i$ and $\phi_{q^\dagger}^j = \phi_{q^\Delta}^i$; $\mathbf{freshhdl}(q)$. Invariants 2, 3, 4, 5 follow.
- Or $w = \{u\}_{k, v}$ with $v \in \mathcal{D}$ we set $\tau_{q^\dagger}^j(\mathbf{freshhdl}(q)) = (u, v)$. For every $h \not\equiv \mathbf{freshhdl}(q)$, put $\tau_{q^\dagger}^j(h) = \tau_{q^\Delta}^i(h)$. Also set $\kappa_{q^\dagger}^j = \kappa_{q^\Delta}^i$ and $\phi_{q^\dagger}^j = \phi_{q^\Delta}^i$; $\mathbf{freshhdl}(q)$. We remark that $\mathbf{comp}_{q^\Delta}(v) = \mathbf{comp}_{q^\Delta}(\mathbf{freshhdl}(q^\Delta))$ from which invariants 2, 3, 4, 5 follow.

Remarking that $\bigvee_{j \in J_i} F_{q^\dagger}^j = F_{q^\Delta}^i$ and for $j \neq j'$ $F_{q^\dagger}^j \wedge F_{q^\dagger}^{j'} = \text{false}$, we get that $F_{q^\dagger}^j$ is a (q_0^\dagger, q^\dagger) -valid partition.

Finally, remark that if we obtain $(F_{q^*}^i)_{i \in I_{q^\dagger}}$, $(\phi_{q^*}^i)_{i \in I_{q^\dagger}}$, $(\kappa_{q^*}^i)_{i \in I_{q^\dagger}}$ and $(\tau_{q^*}^i)_{i \in I_{q^\dagger}}$ by replacing every occurrence of k^\dagger in $(F_{q^\dagger}^i)_{i \in I_{q^\dagger}}$, $(\phi_{q^\dagger}^i)_{i \in I_{q^\dagger}}$, $(\kappa_{q^\dagger}^i)_{i \in I_{q^\dagger}}$ and $(\tau_{q^\dagger}^i)_{i \in I_{q^\dagger}}$ with k^* , as is mandated by invariant 6, then invariants 1, 2, 3, 4 and 5 hold in the $*$ case, since they hold in the \dagger case. \square

We then obtain the proof of the symbolic theorem from the invariants.

Proof of Theorem 2. The main theorem follows easily from the invariants. First we remark that the handle consistency property

$$\phi^\dagger, \neg \mathbf{comp}_{q^\dagger}(h) \rightarrow \bigvee_{h_0 \in \mathcal{H}} \text{EQ}(\mathbf{val}_{q^\dagger}(h), \mathbf{val}_{q_0^\dagger}(h_0)) \sim \phi^\dagger, \text{true}$$

is true if and only if

$$\begin{aligned} [\neg \mathbf{comp}_{q^\dagger}(h)] \bigvee_{k \in K_0} \text{EQ}(\mathbf{val}(h), (k, a_k)) \\ = [\neg \mathbf{comp}_{q^\dagger}(h)] \text{true}. \end{aligned}$$

Let I_\top be the set of indices such that $F_{q^*}^i \rightarrow \neg \mathbf{comp}_{q^*}(h^*) = \text{true}$. We write for $\Delta \in \{*, \dagger\}$:

$$\begin{aligned} F_{q^\Delta}^\perp &\equiv \bigvee_{i \notin I_\top} F_{q^\Delta}^i \\ &= \bigwedge_{i \in I_\top} \neg F_{q^\Delta}^i \text{ (as } (F_{q^\Delta}^i)_{i \in I} \text{ is a valid partition)} \end{aligned}$$

The proof goes as follows: the $(F_{q^*}^i)_{i \in I_\top \cup \{\perp\}}$, $(F_{q^\dagger}^i)_{i \in I_\top \cup \{\perp\}}$ families satisfy the preconditions of the case disjunction axiom, hence it is enough to prove for every i

- 1) The relevant restrictions of the frame are indistinguishable:

$$F_{q^*}^i, [F_{q^*}^i][\neg \mathbf{comp}_{q^*}(h^*)]\phi_i^* \sim F_{q^\dagger}^i, [F_{q^\dagger}^i][\neg \mathbf{comp}_{q^\dagger}(h^*)]\phi_i^\dagger$$

- 2) For every $h, h' \in \mathcal{H}$ the handle consistency property holds:

$$\begin{aligned} \left[F_{q^\dagger}^i \wedge \neg \mathbf{comp}_{q^\dagger}(h)\right] \bigvee_{k \in K_0} \text{EQ}(\mathbf{val}(h), (k, a_k)) \\ = \left[F_{q^\dagger}^i \wedge \neg \mathbf{comp}_{q^\dagger}(h)\right] \text{true} \end{aligned}$$

Let us start by proving 1. First, consider $i \neq \perp$.

By definition of I_\top , we have $F_{q^\dagger}^i \rightarrow \neg \mathbf{comp}_{q^\dagger}(h^*) = \text{true}$. We remark, by invariant 6 that

$$F_{i^*}^*, [F_{i^*}^*]\phi_i^* \equiv \left(F_{i^\dagger}^\dagger, [F_{i^\dagger}^\dagger]\phi_i^\dagger\right) [k^*/k^\dagger]$$

As $F_{i^\dagger}^\dagger$ and ϕ_i^\dagger are $(q_0^\dagger, q^\dagger, F_{q^\dagger}^i)^\dagger$ -normal, k^\dagger only appears under wraps by usable keys with attribute a^* . Let us call K this set of usable keys. Note that the attribute a^* is well used by normality. By the usable secrecy axiom we get

$$F_{i^\dagger}^\dagger, [F_{i^\dagger}^\dagger]\phi_i^\dagger \sim \circ_{K, a^*}(F_{i^\dagger}^\dagger), [\circ_{K, a^*}(F_{i^\dagger}^\dagger)]\circ_{K, a^*}(\phi_i^\dagger).$$

Similarly, we get on the star side

$$F_{i^*}^*, [F_{i^*}^*]\phi_i^* \sim \circ_{K, a^*}(F_{i^*}^*), [\circ_{K, a^*}(F_{i^*}^*)]\circ_{K, a^*}(\phi_i^*).$$

As neither k^* nor k^\dagger appear in K , we still have

$$\begin{aligned} \circ_{K,a^*}(F_i^*), [\circ_{K,a^*}(F_i^*)] \circ_{K,a^*}(\phi_i^*) &\equiv \\ &\left(\circ_{K,a^*}(F_i^\dagger), [\circ_{K,a^*}(F_i^\dagger)] \circ_{K,a^*}(\phi_i^\dagger) \right) [k^*/k^\dagger] \end{aligned}$$

Additionally, note that k^\dagger does not appear in $\circ_{K,a^*}(F_i^\dagger), [\circ_{K,a^*}(F_i^\dagger)] \circ_{K,a^*}(\phi_i^\dagger)$ as all the wraps in which k^\dagger could have appeared are replaced by wraps of fresh keys, hence

$$\begin{aligned} \circ_{K,a^*}(F_i^*), [\circ_{K,a^*}(F_i^*)] \circ_{K,a^*}(\phi_i^*) \\ \equiv \circ_{K,a^*}(F_i^\dagger), [\circ_{K,a^*}(F_i^\dagger)] \circ_{K,a^*}(\phi_i^\dagger) \end{aligned}$$

We conclude by reflexivity and transitivity of \sim .

We now consider the case $i = \perp$. As $(F_{q^\Delta}^i)_{i \in I}$ is a valid partition, by the definition of I_\top we have

$$F_{q^\Delta}^\perp \rightarrow \mathbf{comp}_{q^\Delta}(h^*) = \mathbf{true}$$

We therefore have to prove $F_{q^*}^\perp \sim F_{q^\dagger}^\perp$. We remark that by the definition of normality, the key k^* only appears encrypted under keys that are jointly usable in all $F_{q^*}^i$ for $i \in I_\top$. We conclude

$$\bigvee_{i \in I_\top} F_{q^*}^i \sim \bigvee_{i \in I_\top} F_{q^*}^i$$

by the same arguments we used for the $i \neq \perp$ case. This concludes the proof of point 1.

In order to prove point 2 we start by remarking that if $F_{q^\dagger}^i \rightarrow \mathbf{comp}_{q^\dagger}(h) = \mathbf{true}$ then the property is trivial. We now assume $F_{q^\dagger}^i \rightarrow \neg \mathbf{comp}_{q^\dagger}(h) = \mathbf{true}$. By invariant 3, we have

$$\begin{aligned} [F_{q^\dagger}^i] \mathbf{val}_{q^\dagger}(h) &= [F_{q^\dagger}^i] \kappa_{q^\dagger}^i(h) \\ &= [F_{q^\dagger}^i] (k, a_k) \text{ for some } k \in K_0 \end{aligned}$$

from which point 2 follows immediately. \square

APPENDIX E

PROOF OF THE COMPOSITION THEOREM

In this Section we prove that a key wrapping API using a secure AE-AD wrapping mechanism and a valid policy can be composed with an IND-CCA secure encryption scheme to give a secure wrap and encrypt API.

Proof of Theorem 3. The proof is a fairly standard hybrid argument. We construct four experiments $H_b^{b'}$ for $b, b' \in \{0, 1\}$ which are slight modifications of the experiments $\text{Exp}_b^{\text{COMP}}$. At the start of each new experiment we generate a fresh key k^\dagger . We modify the oracle $\mathcal{O}_b^{\text{enc}}$ so that, if $b' = 0$ and if $h = h^*$, we replace k^* by k^\dagger before encrypting. In all other respects, the experiment $H_b^{b'}$ is identical to $\text{Exp}_b^{\text{COMP}}$. Let $P_{b,b'} = \mathbb{P}[H_b^{b'}(\mathcal{A}, \vec{a}, m) = 1]$. By construction, $H_b^1(\mathcal{A}, \vec{a}, m) = \text{Exp}_b^{\text{COMP}}(\mathcal{A}, \vec{a}, m)$. So we have

$$\begin{aligned} \text{Adv}_{\text{API}'}^{\text{IND-CCA}}(\mathcal{A}) &= \frac{1}{2} (P_{1,1} - P_{0,1}) \\ &= \frac{1}{2} [(P_{1,1} - P_{1,0}) + (P_{1,0} - P_{0,0}) + (P_{0,0} - P_{0,1})]. \end{aligned}$$

We show that there is a key management API satisfying the conditions of Theorem 1 which can be attacked with advantage $\frac{1}{2} (P_{1,1} - P_{1,0})$ or $\frac{1}{2} (P_{0,0} - P_{0,1})$, so these must both be negligible. We also show that there is an IND-CCA adversary against the encryption scheme which wins with probability $\frac{1}{2} (P_{1,0} - P_{0,0})$, so this must also be negligible. Therefore $\text{Adv}_{\text{API}'}^{\text{IND-CCA}}(\mathcal{A})$ is negligible.

Consider the adversary \mathcal{A} in the experiments H_b^1 and H_b^0 . These experiments differ only in whether or not encryption queries using the handle h^* are responded to with encryptions (either of the adversary's message or a random message of the same length) using the real key k^* or a fresh, random key k^\dagger . So we can build an adversary \mathcal{B} in the KEYSEC game for the key wrapping API that uses \mathcal{A} as follows:

First, \mathcal{B} tests h^* at the start of the game, receiving either k^* or a fresh random key k^\dagger . Then, \mathcal{B} submits any other external handles (handles $h_i \neq h^*$ such that $a_i \in \mathcal{E}$) to its `corrupt` oracle, learning the values of these keys. When \mathcal{A} makes any queries to `wrap`, `unwrap` or `corrupt`, \mathcal{B} forwards these to its oracles and returns the responses to \mathcal{A} . For queries to `enc` and `dec` under handles $h \neq h^*$, \mathcal{B} answers faithfully using the corrupted external keys (if the corresponding attributes are not in \mathcal{E} , then \mathcal{B} answers with ζ). For queries to `enc` or `dec` under h^* , \mathcal{B} answers with its test key: encrypting the real message if $b = 1$ and a random message of the same length if $b = 0$. Finally, when \mathcal{A} submits a guess, \mathcal{B} forwards this guess.

It is clear that \mathcal{B} wins the KEYSEC game with advantage $\frac{1}{2} |P_{b,1} - P_{b,0}|$, as long as, according to the sacrifice function, corrupting the external handles other than h^* does not compromise h^* . But from `API` we can build a new API `API*` that also satisfies the conditions of Theorem 1 and in which the corruption of external handles does not compromise the other external handles.

Concretely, we “split” the external attributes from `API` so that each handle in the KEYSEC game for `API*` with an external attribute has a unique attribute. That is, we replace any external attributes that collide among a_1, \dots, a_n in the IND-CCA game played by \mathcal{A} by unique external attributes in the KEYSEC game played by \mathcal{B} . The policy for `API*` simply agrees with the policy for `API` (i.e. the value of the new policy on a new external attribute a is given by the value of the old policy on the external attribute from which a was split). This means that all queries to `API` given by \mathcal{A} that are allowed by the policy in `API` will be allowed by the policy in `API*`. Then we modify the sacrifice function so that, for all external attributes a and a' we have $a \in \text{sacr}(a') \Leftrightarrow a = a'$. This means that the compromise of external attributes will not lead to the compromise of other external attributes (which is a reasonable restriction anyway since one cannot wrap under external keys). Since the policy automatically rejects wraps under external keys, the new policy will be valid with respect to the new sacrifice function. With the new external attributes and new sacrifice function, \mathcal{B} will not compromise h^* in the KEYSEC game for `API*` unless \mathcal{A} does in the IND-CCA game for `API'`. By Theorem 1 applied to `API*`, we have that $P_{b,1} - P_{b,0}$ is negligible for $b \in \{0, 1\}$.

It remains to prove that $P_{1,0} - P_{0,0}$ is negligible. Consider the adversary \mathcal{A} in the experiments H_1^0 and H_0^0 . These experiments differ only in whether encryptions using a key k^\dagger which cannot be wrapped, unwrapped or corrupted are genuine or encryptions of random messages. We construct an adversary \mathcal{B} in the IND – CCA game for the encryption scheme who simulates the IND – CCA API game with its own keys, answering all oracle queries honestly apart from encryption and decryption queries using h^* , which it submits to its IND-CCA oracles. If \mathcal{A} tries to compromise h^* , \mathcal{B} aborts the simulation and outputs a random bit. Otherwise, \mathcal{B} forwards the bit that \mathcal{A} submits. Note that \mathcal{B} is a valid adversary for the IND – CCA encryption game (not decrypting the outputs of the encryption oracle) if \mathcal{A} has non-zero advantage in the API game. Then we have that the advantage $\mathbb{P}_b \left[\mathcal{B}^{\mathcal{O}_{b,k^*}^{\text{enc}}, \mathcal{O}_{k^*}^{\text{dec}}} = b \right] - \frac{1}{2}$ of \mathcal{B} against the IND – CCA game for the encryption scheme is exactly $\frac{1}{2} (P_{1,0} - P_{0,0})$, so this quantity is indeed negligible. \square

APPENDIX F REFINEMENT OF PKCS#11

In this Section, we make precise our secure refinement of the PKCS#11 specification.

Formally, we consider a subset of PKCS#11 operations relating to the management of symmetric keys, namely `C_WrapKey`, `C_UnwrapKey`, `C_GetAttributeValue` and `C_SetAttributeValue`, and show that, if certain attributes that are modifiable in the standard are made unmodifiable, and the wrapping mechanism is a secure AE-AD scheme transmitting the attributes of keys as associated data, then the resulting key wrapping functionality can be securely composed with an IND-CCA secure encryption. Note that the security policy of PKCS#11 must only depend on unmodifiable attributes (otherwise an adversary can simply change the attributes and undermine the policy). Therefore, while modifiable attributes (such as a human-readable name for a key) may aid usability, they are omitted from our discussion.

In PKCS#11, all objects, including keys, consist entirely of a list of attributes. In particular, the cryptographic value of a key is just another object attribute. There is a command, `C_GetAttributeValue`, which returns the list of attributes of an object. However, if the object is a key with `CKA_SENSITIVE` set to `CK_TRUE`, the value of the key will not be revealed. Clearly, any symmetric keys with `CKA_SENSITIVE` set to `CK_FALSE` cannot be used to protect data or other keys. Therefore, we assume that all keys have been generated on the token prior to its execution and that all keys have the attribute `CKA_SENSITIVE` set to `CK_TRUE`.

While `C_SetAttributeValue` can be used to modify certain attributes, it cannot be used to change `CKA_SENSITIVE` from `CK_TRUE` to `CK_FALSE` (this attribute is *sticky-on*) or `CKA_EXTRACTABLE` from `CK_FALSE` to `CK_TRUE` (this attribute is *sticky-off*). However, by default, the attributes `CKA_WRAP` and `CKA_UNWRAP` are modifiable, so it is impossible to have a set of *external* keys that cannot be used for wrapping (and therefore remain indistinguishable

from random) without strengthening the PKCS#11 attribute policy. Therefore we impose the following additional restriction, which is both for security and convenience: the attributes `CKA_WRAP`, `CKA_UNWRAP`, `CKA_EXTRACTABLE`, `CKA_TRUSTED` and `CKA_WRAP_WITH_TRUSTED` are both sticky-on and sticky-off (i.e. unmodifiable).

The attribute `CKA_TRUSTED` is used by PKCS#11 to identify *trusted wrapping keys* and only has a meaning for those keys such that `CKA_WRAP = CK_TRUE`. On the other hand, the attribute `CKA_WRAP_WITH_TRUSTED` identifies keys that may only be *wrapped under* trusted wrapping keys. This attribute only has a meaning for those keys such that `CKA_EXTRACTABLE = CK_TRUE`.

A policy in our model specifies what keys a particular key may wrap. Considering the attributes `CKA_WRAP`, `CKA_EXTRACTABLE`, `CKA_TRUSTED` and `CKA_WRAP_WITH_TRUSTED`, the PKCS#11 documentation identifies nine meaningful *key templates* (choices of values for certain attributes). We denote these key templates as follows: a_x^y is the template of a key that

- cannot wrap if $x = 0$ and is an untrusted (resp. trusted) wrapping key if $x = u$ (resp. $x = t$)
- cannot be wrapped if $y = 0$ and can be wrapped under trusted (resp. trusted or untrusted) keys if $y = t$ (resp. $y = u$).

We obtain a secure subset of the wrapping operations permitted by PKCS#11 as follows:

Let $\mathcal{E} = \{a_0^0, a_0^t, a_0^u\}$ and let $\mathcal{D} = \mathcal{E} \cup \{a_u^0, a_u^t, a_t^0\}$. We note that the PKCS#11 policy, viewed as a function on $\mathcal{D} \times \mathcal{D}$, prevents wrapping under external keys as required. Furthermore, \mathcal{D} is maximal: if there were keys of templates a_t^t , a_t^u or a_u^u , then the policy graph would have a cycle, which is forbidden in our model. Thus, by Theorem 1, for any sacrifice function on templates respected by the policy (i.e. any graph on the six templates with the wrapping policy as a subgraph), any external keys not trivially compromised according to the sacrifice function will be indistinguishable from random in any execution of the key wrapping API, provided that some conditions on the wrapping mechanism are satisfied. Moreover, Theorem 3 shows that the external keys can be used by an IND-CCA secure encryption scheme and the wrapping mechanism will not interfere with the security of the encryption scheme.

We note that imposing additional restrictions on attribute combinations as we have done is explicitly authorised by the PKCS#11 documentation. On Page 131, it reads: “to partition the wrapping keys so they can only wrap a subset of extractable keys the attribute `CKA_WRAP_TEMPLATE` can be used on the wrapping key to specify an attribute set that will be compared against the attributes of the key to be wrapped” [OAS15a].

Our security theorem for key wrapping APIs requires that the wrapping mechanism is a secure AE-AD scheme that automatically sends the attributes of keys as associated data when wrapping and chooses a fresh, random nonce for each wrap. This is explicitly supported in PKCS#11, since version

2.4 of the standard added support for the use of AES-GCM for key wrapping with arbitrary IVs (nonces) and arbitrary associated data [OAS15b]. One could easily configure the wrap command to generate a random IV, retrieve the associated data using `C_GetAttributeValue` and output these parameters along with the ciphertext.

We remark that the template policy we consider here is more general than the one proven correct in [Kün15], at the cost of assuming that the wrapping scheme is AE-AD secure (or possibly a secure Deterministic Authenticated Encryption scheme) and that attributes are transmitted properly through the associated data. In particular, the policy proposed in [Kün15] prevents transmitting any wrapping key between devices, while ours allows this behaviour for the template a_u^t (an untrusted wrapping key that can be wrapped under any trusted wrapping key). This behaviour allows devices to exchange one long term managing key, and then use it to establish short-lived external keys. This would typically be needed for creating session keys.