A Taxonomy of Cryptographic Irreversibility: The Bateson Framework and a Novel Hybrid Sponge-LWE Construction

Kevin Fathi*

Abstract

We introduce the Bateson Framework, a conceptual model that provides a taxonomy for computational irreversibility based on two fundamental sources: information loss (degeneracy) and injected uncertainty (ambiguity). This model allows for a clear classification of existing primitives, such as hash functions (high-degeneracy) and LWE-based schemes (high-ambiguity). Our primary contribution is to demonstrate that this taxonomy is not merely descriptive but generative. We present illustrative instantiations based on the Learning with Errors (LWE) assumption and cryptographic sponges, analyzing their security and practical limitations. Crucially, we introduce a novel hybrid construction that simultaneously leverages both degeneracy and ambiguity by composing a cryptographic sponge with an LWE scheme, showcasing the power of the unified approach for designing robust cryptographic primitives.

Keywords: One-Way Functions, Information Theory, Learning with Errors (LWE), Cryptographic Hash Functions, Taxonomy, Hybrid Cryptography, Post-Quantum Cryptography.

1 Introduction

The concept of a one-way function (OWF) is the cornerstone of modern cryptography.[8] While the existence of OWFs remains a conjecture (implying $P \neq NP$), candidate constructions are typically based on the presumed hardness of specific computational problems, such as integer factorization or structured problems in lattices.[12] This paper approaches the construction of one-wayness from a different perspective, rooted in information theory and formal models of keyed functions.

We introduce the Bateson Framework, inspired by Gregory Bateson's concept that information is context-dependent.[1] This framework models a cryptographic function where a public input (message) is processed under a secret context (frame, e.g., a key), and this frame is suppressed in the output. The foundation of this approach draws conceptual inspiration from the Generalized Entropy-Complexity Correspondence (GECC) [4], which suggests that irreversibility can arise from two distinct sources:

1. **Degeneracy**: Information loss due to many-to-one mappings (e.g., compression in hash functions).
2. **Ambiguity**: Uncertainty due to the injection of noise in probabilistic functions (e.g., the error term in LWE) or the use of a large secret space.

The Bateson Framework unifies these two perspectives, providing a generalized taxonomy for understanding and, crucially, constructing cryptographic irreversibility.

*Independent Researcher, fathikevin@protonmail.com

1.1 Our Contributions

This paper's primary contribution is to introduce this novel taxonomy and demonstrate its generative power by constructing a novel hybrid one-way function. Our main contributions are:

1. **An Information-Theoretic Taxonomy:** We formalize a taxonomy that distinguishes and unifies irreversibility stemming from information loss (Degeneracy, related to $H(Y|X)$) and injected noise (Ambiguity, related to $H(X|Y)$).
2. **A Novel Hybrid Construction:** We introduce a new one-way function design that simultaneously leverages both degeneracy (via a cryptographic sponge) and ambiguity (via LWE). This demonstrates the generative power of the framework and presents a new design pattern.
3. **Illustrative Instantiations and Analysis:** We provide two theoretical instantiations to demonstrate the framework's breadth: a high-ambiguity construction based on standard LWE, and a high-degeneracy construction using cryptographic sponges, along with their security analyses and practical limitations.

2 Preliminaries

2.1 Notation

We denote the security parameter by n. A function $\epsilon(n)$ is negligible, denoted $negl(n)$, if it vanishes faster than the inverse of any polynomial in n. PPT stands for Probabilistic Polynomial Time. We use $x \leftarrow \mathcal{D}$ to denote sampling x from a distribution $\mathcal{D}$, and $x \leftarrow U(S)$ for uniform sampling from a set S.

2.2 Cryptographic Definitions

Definition 1. *A function $f : \{0,1\}^* \to \{0,1\}^*$ (potentially probabilistic) is a strong one-way function if it is easy to compute (in PPT) and hard to invert. Hard to invert means for every non-uniform PPT adversary $\mathcal{A}$:*

$$\Pr[x \leftarrow \{0,1\}^n; y \leftarrow f(x); x' \leftarrow \mathcal{A}(1^n, y) : y \in supp(f(x'))] \leq negl(n). \tag{1}$$

2.3 Learning with Errors (LWE)

Definition 2. *For parameters n (dimension), m (number of samples), q (modulus), and an error distribution χ over $\mathbb{Z}_q$, the search-LWE problem requires finding the secret $s \in \mathbb{Z}_q^n$ given $(A, b = sA + e \pmod q)$, where $A \leftarrow U(\mathbb{Z}_q^{n \times m})$, $s \leftarrow U(\mathbb{Z}_q^n)$, and $e \leftarrow \chi^m$.*

The hardness of LWE is based on the conjectured hardness of worst-case lattice problems, making it a leading candidate for post-quantum cryptography.[12]

3 The General Bateson Framework

3.1 The Generalized Bateson Function

The Bateson function formalizes the process where an input is interpreted under a context, and that context is hidden in the output.

Definition 3. *A family $\{G_n\}_{n\in\mathbb{N}}$ of Probabilistic Bateson Grammars is a collection of tuples $G_n = (M_n, F_n, \Theta_n, P_{G_n})$ where:*

- *M_n is the message space (public input).*
- *F_n is the set of interpretive frames (secret context).*
- *Θ_n is the output space.*
- *$P_{G_n}(y|m, f)$ is an efficiently samplable Probabilistic Canonicalization Function, defining the probability of output y given message m and frame f.*

Definition 4. *Given $\{G_n\}$, the Bateson function $f^{+}_{B,n} : M_n \times F_n \to M_n \times \Theta_n$ is a probabilistic function where, on input (m, f), the output is (m, y), where y is sampled according to $P_{G_n}(y|m, f)$.*

The crucial characteristic is **Frame Suppression**: the output includes the message m but suppresses the frame f.

3.2 A Taxonomy of Irreversibility

The GECC identity, $H(Y) = H(X) + H(Y|X) - H(X|Y)$, provides a conceptual lens for our framework.[4] However, for rigorous cryptographic proofs, we must move beyond Shannon entropy, which measures average uncertainty, to measures that provide worst-case guarantees.[8]

- **Min-entropy (H_∞)** measures the unpredictability of a random variable by focusing on its most likely outcome. This is crucial for arguments involving randomness extractors like the Leftover Hash Lemma (LHL).
- **Statistical Distance** ($\Delta(P, Q)$) measures the indistinguishability of two probability distributions. A function exhibits high ambiguity if its output distribution is statistically close to uniform.

This perspective allows for a classification of cryptographic primitives based on whether their irreversibility primarily stems from Degeneracy (information loss) or Ambiguity (injected noise/uncertainty).

Limitations of the Taxonomy. It is crucial to note the limitations of this taxonomy. While it effectively captures irreversibility stemming from information loss (degeneracy) and injected noise (ambiguity), it does not encompass all sources of cryptographic one-wayness. As illustrated by the entry for textbook RSA in Table 1, trapdoor permutations rely purely on computational hardness (e.g., the difficulty of integer factorization). Textbook RSA is a deterministic permutation, exhibiting Zero Degeneracy (it is one-to-one) and Zero Ambiguity (no noise is injected). The Bateson Framework therefore provides a unifying view of hash functions and noisy lattice schemes, but it is not exhaustive of all cryptographic constructions.

4 Illustrative Instantiations

We now present two instantiations that illustrate the extreme ends of the taxonomy: a high-ambiguity function based on LWE, and a high-degeneracy function based on cryptographic sponges.

Table 1: A Taxonomy of Primitives in the Bateson Irreversibility Space

Primitive	**Primary Source of Irreversibility**	**Degeneracy Level**	**Ambiguity Level**	**Analogy ($(m, f) \to y$)**
SHA-3	Degeneracy (Compression)	High	Zero	$f(m, f_{\text{fixed}}) = \text{Sponge}(m)$
HMAC	Degeneracy (Keyed Compression)	High	Zero	$f(m, k) = \text{HMAC}(m, k)$
LWE-PKE	Ambiguity (Noise)	Low	High	$f(A, s) = As + e$
RSA (Textbook)	Computational Hardness*	Zero	Zero	$f(m, (N, e)) = m^e \bmod N$
OTP	Ambiguity (Key)	Zero	High	$f(m, k) = m \oplus k$

4.1 Instantiation 1: High Ambiguity via LWE

We present a theoretical instantiation based on the standard LWE assumption. This serves as a canonical example of a high-ambiguity function.

Definition 5. *Let parameters be n, m, q, χ.*

- ***Frame (f):*** *A secret vector $s \in \mathbb{Z}_q^n$.*
- ***Message (m):*** *A public matrix $A \in \mathbb{Z}_q^{n \times m}$.*
- ***Canonicalization (P_G):*** *Computes $y = sA + e \pmod q$, where $e \leftarrow \chi^m$.*

The function is $f_{LWE}(A, s) = (A, sA + e)$.

4.1.1 Security Analysis

The security of this construction relies on the injectivity of the mapping from the secret components (frame and noise) to the output, given the public message.

Lemma 1 (LWE Injectivity)**.** *For typical LWE parameters (e.g., $m > n \log q + \omega(\log n)$, q prime, $q \gg B$ where B is the error bound of χ), the mapping $g_A(s, e) = sA + e$ is injective with overwhelming probability over the choice of A.*

Proof. Assume a collision: $s_1A + e_1 = s_2A + e_2$ with $(s_1, e_1) \neq (s_2, e_2)$. If $s_1 \neq s_2$, let $\Delta s = s_1 - s_2 \neq 0$. Then $\Delta sA = e_2 - e_1 = \Delta e \pmod q$. The vector Δe is "short," as its coefficients are bounded by $2B$.

By the Leftover Hash Lemma (LHL), the family of functions $\{h_A : \mathbb{Z}_q^n \to \mathbb{Z}_q^m\}_{A \in \mathbb{Z}_q^{n \times m}}$ defined by $h_A(x) = xA$ is a universal hash family. For any fixed $\Delta s \neq 0$, the distribution of ΔsA is statistically close to the uniform distribution over $\mathbb{Z}_q^m$. The statistical distance Δ is bounded by $\Delta(\Delta sA, U(\mathbb{Z}_q^m)) \leq \frac{1}{2}\sqrt{2^{m \log q}/2^{H_\infty(\Delta s)}}$.

The probability that this near-uniform vector falls into the small ball of radius $2B$ around the origin is negligible. The number of vectors in this ball is approximately $(4B+1)^m$. The probability of a collision for a fixed Δs is thus bounded by $\frac{(4B+1)^m}{q^m} + \Delta$, which is negligible for $q \gg B$.

By applying a union bound over all possible $\Delta s \neq 0$, the probability that there exists any Δs such that ΔsA is short is negligible. Thus, with overwhelming probability over the choice of A, we must have $s_1 = s_2$. This implies $e_1 = e_2$, contradicting our assumption that $(s_1, e_1) \neq (s_2, e_2)$. □

Theorem 1. *If search-LWE is hard for (n, m, q, χ), then f_{LWE} is a secure one-way function.*

Proof. Assume a PPT adversary $\mathcal{A}$ inverts f_{LWE} with non-negligible probability $\epsilon(n)$. We construct an algorithm $\mathcal{B}$ to solve search-LWE.

Setup: $\mathcal{B}$ receives an LWE instance $(A^*, b^* = s^*A^* + e^*)$.

Reduction: $\mathcal{B}$ invokes $\mathcal{A}$ on (A^*, b^*). If $\mathcal{A}$ returns a preimage (A', s'), $\mathcal{B}$ outputs s'.

Analysis: A valid preimage must satisfy $A' = A^*$ and $s'A^* + e' = b^*$ for some $e' \leftarrow \chi^m$. We have: $s'A^* + e' = s^*A^* + e^*$.

Let E_{Inj} be the event that g_{A^*} is injective (as per Lemma 1). We have $\Pr[E_{Inj}] \geq 1 - negl(n)$. Conditioned on E_{Inj}, the equality implies $(s', e') = (s^*, e^*)$, so $s' = s^*$. The success probability of $\mathcal{B}$ is at least $\epsilon(n) - \Pr[\neg E_{Inj}] = \epsilon(n) - negl(n)$, which is non-negligible. This contradicts the LWE hardness assumption.

The reduction is tight up to the negligible probability of injectivity failure. For standard post-quantum parameters, the concrete bounds from the LHL make this failure probability cryptographically negligible (e.g., $< 2^{-128}$). □

4.1.2 Practical Considerations and Limitations

While this LWE instantiation is provably secure under the standard LWE assumption, its practical viability as a standalone primitive is severely limited. This limitation stems from the requirements of the standard LWE assumption used in Theorem 1.

The security proof requires the public matrix A (the message m) to be uniformly random and unstructured. For NIST Level 1 parameters (e.g., $n = 512, q = 3329$), an uncompressed matrix $A \in \mathbb{Z}_q^{n \times n}$ would require storing $n^2 = 512^2 = 262{,}144$ coefficients. With each coefficient requiring $\lceil \log_2(3329) \rceil = 12$ bits, the total size is $262{,}144 \times 12/8 = 393{,}216$ bytes (approximately 393 KB).

This contrasts sharply with modern post-quantum standards like CRYSTALS-Kyber.[10] Kyber is based on the Module-LWE (MLWE) assumption, which allows the public matrix to have algebraic structure. This structure enables the matrix to be generated from a short seed (e.g., 800 bytes for Kyber-512), dramatically reducing the communication overhead.

Therefore, this LWE instantiation should be viewed as a theoretical construction that illustrates the concept of high-ambiguity irreversibility within the taxonomy, rather than a practical competitor to efficient structured-lattice schemes.

4.2 Instantiation 2: Maximal Degeneracy via Sponge Construction

This construction is functionally equivalent to a standard keyed sponge, such as that used in KMAC.[2] It illustrates a pure degeneracy-based primitive within the taxonomy.

Definition 6. *Let $\pi : \{0,1\}^b \to \{0,1\}^b$ be a public permutation (capacity c, rate r, $b = c + r$). Output length L.*

Table 2: Comparative Analysis: Theoretical Standard LWE-OWF vs. Practical ML-KEM (Kyber)

Feature	**Standard LWE-Based OWF (Theoretical)**	**CRYSTALS-Kyber (Practical)**
Primitive Type	One-Way Function (OWF)	Key Encapsulation Mechanism (KEM)
Hardness Assumption	Standard LWE (unstructured matrix)	Module-LWE (structured matrix)
Public Input Size (NIST L1)	Matrix A ($\approx$393 KB)	Seed for matrix $\mathbf{A}$ (800 bytes)
Role	Illustration of Ambiguity	Practical Key Exchange

- ***Frame*** (f)***:*** *Secret key* $k \in \{0,1\}^c$.
- ***Message*** (m)***:*** *Input data* $m \in \{0,1\}^*$.
- ***Canonicalization*** (P_G)***:*** *Deterministic. The state is initialized with the key (e.g.,* $S_0 = k \,\|\, 0^r$*). m is absorbed via π. Output y is squeezed.*

4.2.1 Security Analysis

Lemma 2 (Sponge Injectivity)**.** *In the Ideal Permutation Model (IPM), the mapping* $F_m(k) = Sponge(\pi, m, k)$ *is injective with respect to the frame k, with overwhelming probability, provided L and c are $\Omega(n)$.*

Proof. The sponge construction is indifferentiable from a Random Oracle (RO).[3] The collision probability for two different keys k_1, k_2 is bounded by $\Pr[F_m(k_1) = F_m(k_2)] \leq \frac{1}{2^L} + O\left(\frac{N^2}{2^c}\right)$, where N is the number of queries. For $L, c = \Omega(n)$, this is negligible, implying injectivity with overwhelming probability. □

Theorem 2. *In the IPM, the Sponge-Based function is a secure one-way function.*

Proof. Inversion requires finding a preimage for a keyed sponge. Standard security arguments for sponge constructions show this requires $\Omega(2^{\min(L,c/2)})$ queries to the ideal permutation. If $c, L = \Omega(n)$, this is computationally infeasible. □

5 The Hybrid Construction: Leveraging the Unified Framework

The true generative power of the Bateson Framework lies in its ability to inspire novel constructions that do not rely solely on one source of irreversibility. By unifying degeneracy and ambiguity under a single taxonomy, we can design hybrid primitives that combine the strengths of both. We introduce a novel hybrid construction that composes a cryptographic sponge (degeneracy) with an LWE scheme (ambiguity).

Definition 7 (Hybrid Sponge-LWE Function)**.**

- ***Message*** (m)***:*** *Arbitrary length bit strings* $\{0,1\}^*$.

- ***Frame (f):*** *A secret master key $k \in \{0,1\}^c$.* ***Output (Θ):*** *LWE ciphertexts $y \in \mathbb{Z}_q^m$.*
Canonicalization ($P_G(y|m,k)$):

1. ***Parameter Derivation (Degeneracy):*** *Initialize a cryptographic sponge (e.g., SHAKE) with the frame k. Absorb the message m.*
2. ***Squeezing:*** *Squeeze the sponge state to derive the LWE secret vector $s \in \mathbb{Z}_q^n$ and a public seed, $seed_A$.*
3. ***Matrix Generation:*** *Expand $seed_A$ using an XOF to generate the public LWE matrix $A \in \mathbb{Z}_q^{n \times m}$.*
4. ***Noise Injection (Ambiguity):*** *Sample a short error vector $e \leftarrow \chi^m$.*
5. ***Output Computation:*** *Compute the output $y = As + e \pmod q$.*

5.1 Security Rationale

This construction's one-wayness stems from two complementary and reinforcing sources, directly corresponding to the taxonomy:

- **Degeneracy (Sponge Component):** The sponge compresses the variable-length input (m, k) into a fixed-size state to derive the intermediate LWE secret s. An adversary attempting to invert the function must first contend with the difficulty of reversing this compression. This requires mounting a preimage attack on the keyed sponge to find an input (m', k') that yields the internal secrets $(s, seed_A)$ consistent with the output y.
- **Ambiguity (LWE Component):** The LWE error term e masks the relationship between the derived secret s and the output y. The adversary faces solving an LWE instance derived from the sponge output.

This hybrid design offers potential advantages over constructions relying on a single hardness assumption. The security appears to rely on a composite assumption: the hardness of preimage attacks on the keyed sponge (often analyzed in the IPM) and the hardness of the LWE problem. This composition represents a genuinely new design pattern facilitated by the unified perspective of the Bateson Framework. A full, rigorous security analysis of this hybrid function is the primary direction for future work.

6 Implementation and Physical-World Context

6.1 Implementation Security and Side-Channel Resistance

Both the Sponge and LWE components utilized in the instantiations, particularly the hybrid construction, are vulnerable to side-channel analysis (SCA) if implemented naively.[9]

- **LWE Component:** The computation $y = sA+e$ is vulnerable to Differential Power Analysis (DPA). Furthermore, the error sampling must be implemented in constant time to prevent timing leakages that could reveal e. Standard countermeasures include algorithmic **masking** (splitting secrets into multiple shares) and **shuffling** the order of operations.
- **Sponge Component:** Power analysis on the internal state during absorption can reveal the secret frame k.

7 Future Work and Conclusion

7.1 Future Work

The framework and the resulting hybrid construction open several critical avenues for future research:

1. **Rigorous Security Analysis of the Hybrid Instantiation:** The most important next step is to provide a full, rigorous security proof for the hybrid Sponge-LWE function (Definition 7) under a composite assumption (e.g., IPM + LWE hardness).

2. **Optimization using Structured Lattices:** Investigate adapting the hybrid construction to use Module-LWE or Ring-LWE to achieve practical efficiency comparable to NIST PQC standards.

3. **Exploring Other Hybrid Instantiations:** Model other post-quantum assumptions, such as those from code-based or isogeny-based cryptography, within the taxonomy and explore new hybrid designs.

7.2 Conclusion

We have introduced the Bateson Framework, providing a novel taxonomy that unifies irreversibility arising from information loss (degeneracy) and injected noise (ambiguity). We analyzed the limitations of this taxonomy, noting it does not encompass purely computational hardness such as trapdoor permutations. We demonstrated the utility of this taxonomy by classifying existing primitives and illustrating the concepts with theoretical instantiations based on LWE and sponges, highlighting the practical limitations of relying on unstructured LWE.

Crucially, we showed that the framework is generative, leading to the design of a novel hybrid construction that simultaneously leverages both degeneracy and ambiguity. By unifying different sources of one-wayness under a single information-theoretic umbrella, this taxonomy offers a powerful new perspective for the design and analysis of cryptographic primitives.

References

[1] Bateson, G.: Steps to an Ecology of Mind. Chandler Publishing Company (1972)

[2] Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Sponge Functions. In: Ecrypt Hash Workshop (2007)

[3] Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: On the Indifferentiability of the Sponge Construction. In: EUROCRYPT. pp. 181-197. Springer (2008)

[4] Fathi, K.: Bridging Shannon and Turing: A Formal Entropy-Complexity Correspondence over Symbolic Structures (2025)

[5] Fathi, K.: Bridging Shannon and Turing in Non-Ideal Systems: Symbolic Entropy in Ambiguous and Uncertain Environments (2025)

[6] Goldreich, O.: Foundations of Cryptography: Volume 1, Basic Tools. Cambridge University Press (2001)

[7] Howe, J., Guneysu, T., Regazzoni, F.: Side-Channel Analysis of Post-Quantum Cryptography: A Review. ACM Computing Surveys (2023)

[8] National Institute of Standards and Technology: Selected Algorithms for Post-Quantum Cryptography (2022)

[9] Pappu, R., Recht, B., Taylor, J., Gershenfeld, N.: Physical One-Way Functions. Science 297(5589), 2026-2030 (2002)

[10] Regev, O.: On Lattices, Learning with Errors, Random Linear Codes, and Cryptography. In: STOC. pp. 84-93. ACM (2005)