

The Entropy Ratio: Quantitative Confidentiality for Trapdoor Computing

Alexander Towell
lex@metafunctor.com

April 2026

Abstract

A cipher map is a total function on bit strings implementing a trapdoor approximation of a latent function: the untrusted machine evaluates it blindly; only the trusted machine, holding the secret, can encode inputs and decode outputs. We develop a quantitative confidentiality theory for cipher map systems. The central measure is the *entropy ratio* $e = H/H^*$, a normalized form of Shannon leakage from the Quantitative Information Flow literature. Our contribution is a bridge from the cipher map framework to this measure: the representation-uniformity parameter δ controls the ratio via the Fannes–Audenaert continuity inequality, giving $e \geq 1 - \delta - h_2(\delta)/n$, which reduces confidentiality engineering to the concrete problem of minimizing δ . We analyze two constructions that reduce δ with explicit costs (noise injection and multiple representations), and inherit a third, encoding granularity, from the cipher maps framework. We prove a compositional leakage theorem: even when marginals are δ -uniform, the joint distribution over shared cipher values is recoverable at standard parametric rate, showing that reducing δ is necessary but not sufficient. Experimental validation on encrypted Boolean search over the 20 Newsgroups corpus confirms the theoretical predictions for FPR compounding and the encoding granularity spectrum.

1 Introduction

Given a cipher map system—a trusted machine encoding queries, an untrusted machine evaluating total functions on bit strings, and results decoded only by the trusted machine—how much does the untrusted machine learn?

The cipher map framework [14] provides four properties (totality, representation uniformity, correctness, composability) parameterized by $(\eta, \varepsilon, \delta, \mu)$ that characterize what the untrusted machine can and cannot do. Totality ensures every input produces output; representation uniformity bounds frequency analysis; correctness governs error rates; composability enables chained evaluation. But these properties are *qualitative* guarantees. A system designer needs *quantitative* answers: given a specific deployment with a known query distribution, known vocabulary size, and known resource budget, how confidential is the system, and how can it be improved?

This paper develops the quantitative theory. The central idea is the *entropy ratio* $e = H/H^*$: the ratio of observed entropy in the cipher map output stream to the maximum entropy achievable under system constraints. The measure itself is not new; it is a normalized form of Shannon leakage studied in the Quantitative Information Flow literature (§2). The contribution here is the connection to the cipher map framework: we prove that the representation-uniformity parameter δ directly lower-bounds the entropy ratio via the Fannes–Audenaert continuity inequality [6, 2],

$$e \geq 1 - \delta - h_2(\delta)/n,$$

where h_2 is the binary entropy. This reduces confidentiality engineering to the concrete problem of minimizing δ .

We analyze two constructions that reduce δ with explicit costs:

1. **Noise injection.** The trusted machine mixes R filler queries per N real queries. By totality, the untrusted machine cannot distinguish real from filler. Cost: bandwidth $(1 + R/N)$. Effect: Fisher information for any correlation estimator is reduced by $\rho^2 = (N/(N + R))^2$.
2. **Multiple representations.** Each element receives $K(x) \geq 1$ encodings, with $K(x) \propto D(x)$ to flatten the cipher value distribution (classical homophonic substitution: frequent elements get more code symbols so that per-cipher frequency equalizes). Cost: space $(\sum_x K(x)$ slots). Effect: $\delta \rightarrow 0$ as the normalizing constant grows.

A third construction, **encoding granularity** (controlled by an entanglement parameter p that interpolates between component-wise encoding with correlation leakage and joint encoding with space $O(|Y|^k)$), is inherited from the cipher maps framework [14, Sec. 9]. We use it but do not re-prove it here.

The paper’s compositional contribution is distinct from either construction. We prove (Theorem 6.1) that δ -uniform marginals do not prevent joint recovery when the untrusted machine observes multiple evaluations on a shared cipher value: the joint distribution is recoverable at standard parametric rate $O(|Y_1| \cdot |Y_2|/\xi^2)$. This is the reason practitioners cannot stop at multiplicity. We also include, for completeness, FPR compounding through Boolean chains, which is inherited from [14, Sec. 8].

Finally, we provide practical measurement tools. The entropy of a cipher map output stream can be estimated by compressing it: more compressible means more predictable means less confidential. This avoids the need for explicit probabilistic models of the adversary.

Contributions.

1. **Fannes bridge.** The representation-uniformity parameter δ directly lower-bounds the entropy ratio, $e \geq 1 - \delta - h_2(\delta)/n$, via the Fannes–Audenaert continuity inequality (Theorem 4.1). This makes δ the operational handle for confidentiality engineering in cipher map systems (§4).
2. **Two constructions that reduce δ .** Noise injection (bandwidth cost $1 + R/N$, Fisher-information dilution ρ^2 ; Theorem 5.1) and multiple representations (space cost $\sum_x K(x)$ with $K(x) \propto D(x)$, the classical homophonic prescription; Theorem 5.2). A third construction, encoding granularity, is inherited from the cipher maps framework and used without re-proof (§5).
3. **Compositional leakage bound.** Theorem 6.1 shows that δ -uniform marginals do not prevent joint recovery when the untrusted machine observes multiple evaluations on a shared cipher value: the joint distribution is recoverable from $O(|Y_1| \cdot |Y_2|/\xi^2)$ observations. This is the paper’s main compositional contribution and is the reason reducing δ is necessary but not sufficient. FPR compounding through Boolean chains is included for context; it is inherited from [14] (§6).
4. **Practical measurement and validation.** Compression-based entropy estimation avoiding explicit adversary models, and a case study demonstrating confidentiality improvement from 72% to 98% with approximately $1.04\times$ space and $1.5\times$ bandwidth overhead (§7, §8).

What this paper does not do. We do not use ORAM, FHE, or simulation-based security definitions. Privacy in the cipher map model comes from the one-way hash, the totality of \hat{f} , and representation uniformity—not from access-pattern indistinguishability or computational hardness assumptions beyond the hash function. The framework is information-theoretic throughout.

2 Related Work

Quantitative information flow. The entropy ratio $e = H/H^*$ is a normalized form of Shannon leakage, studied extensively in the QIF literature [12, 1]. Smith [12] proposes min-entropy leakage as a worst-case measure; Alvim et al. [1] provide a comprehensive treatment of information-theoretic leakage measures. The measure itself is not our contribution, and we do not claim novelty there. What is new here is (i) the *Fannes bridge* connecting the representation-uniformity parameter δ of cipher maps to the entropy ratio ($e \geq 1 - \delta - h_2(\delta)/n$), making δ the operational handle QIF does not have, and (ii) the compositional leakage bound showing that δ -uniform marginals are insufficient when shared variables recur across evaluations (Theorem 6.1). The QIF literature treats leakage of a function; the cipher map setting adds a four-parameter construction $(\eta, \varepsilon, \delta, \mu)$ and a composition discipline under which that leakage must be reasoned about. We discuss the Shannon-vs-min-entropy choice in §9.

SSE leakage mitigation. Noise injection for encrypted search has been studied in the SSE literature. Bost and Fouque [3] propose dummy queries with game-based security analysis; Demertzis et al. [5] develop tunable leakage-functionality trade-offs. Our analysis differs in using information-theoretic measures (entropy ratio) rather than simulation-based definitions, and applies to general cipher maps, not only encrypted search.

Leakage from property-preserving encryption. Naveed et al. [9] and Islam et al. [7] demonstrate inference attacks exploiting preserved algebraic structure. The cipher map model avoids property preservation entirely: the untrusted machine evaluates a total function on opaque bit strings and cannot determine the output type from the cipher value.

3 Preliminaries

We recall the cipher map abstraction from [14]. All notation follows that paper; we cite rather than re-derive.

Definition 3.1 (Cipher map [14, Def. 3.1]). A *cipher map* for a latent function $f : X \rightarrow Y$ is a tuple $(\hat{f}, \text{enc}, \text{dec}, s)$ where $\hat{f} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a total function on n -bit strings, $\text{enc} : X \times \{0, \dots, K(x)-1\} \rightarrow \{0, 1\}^n$ is an encoding function with $K(x) \geq 1$ representations per element, $\text{dec} : \{0, 1\}^n \rightarrow Y \cup \{\perp\}$ is a decoding function, and s is a secret (the trapdoor). The untrusted machine holds \hat{f} . The trusted machine holds enc , dec , and s .

The four properties, parameterized by $(\eta, \varepsilon, \mu, \delta)$ [14, Sec. 4]:

Totality. \hat{f} is defined on all 2^n inputs. Out-of-domain outputs are indistinguishable from uniform under the random oracle model.

Representation uniformity (δ -bounded). The marginal distribution of cipher values is δ -close to uniform over $\{0, 1\}^n$ in total variation distance: $d_{\text{TV}}(P, U) = \frac{1}{2} \sum_c |P(c) - U(c)| \leq \delta$.

Informally, δ bounds the advantage of any test that tries to distinguish real cipher values from random ones.

Correctness (η -bounded). The fraction of domain elements that decode incorrectly is at most η :

$$\Pr[\text{dec}(\hat{f}(\text{enc}(x, k))) \neq f(x)] \leq \eta.$$

Nonzero η provides plausible deniability (an observed result might be wrong) and reduces construction cost.

Noise-decode probability (ε). A random n -bit string decodes to a valid output with probability ε . Lower ε means out-of-domain queries are more likely to be rejected (decoded as \perp), at the cost of more bits per element.

Value cost ($\mu = H(Y)$). The average number of bits needed to encode one output value, equal to the Shannon entropy of the output distribution.

Composability. The composition of two cipher maps is again a cipher map: $\hat{g} \circ \hat{f}$ has correctness $\eta_{g \circ f} \leq 1 - (1 - \eta_f)(1 - \eta_g)$. Errors accumulate predictably, enabling pipeline analysis.

The space per element decomposes as bits/element = $-\log_2 \varepsilon + H(Y)$, where $-\log_2 \varepsilon$ bits distinguish signal from noise and $H(Y) = \mu$ bits encode the function value [14, Thm. 6.1].

Trusted/untrusted model. The trusted machine T encodes inputs, decodes results, and injects filler queries. The untrusted machine U evaluates \hat{f} on any bit string and returns results. U cannot decode, distinguish real from filler queries, determine the domain X , or enumerate which inputs are “in-domain” [14, Sec. 5].

Acceptance predicates. The batch construction of a cipher map is unified under an *acceptance predicate*: a family of disjoint sets $\{A(y) \subseteq \{0, 1\}^n : y \in Y\}$ that partitions hash space among output values [14, Sec. 6.2]. Shannon-optimal allocation sets $|A(y)| \propto \Pr[f(X) = y]$, simultaneously minimizing space and maximizing output indistinguishability.

Information-theoretic notation. We use standard information-theoretic quantities throughout. *Shannon entropy* $H(X) = -\sum_x p(x) \log_2 p(x)$ measures the average surprise (in bits) of a random variable; higher entropy means less predictable. *Conditional entropy* $H(X | Y)$ is the remaining uncertainty in X after observing Y . *Mutual information* $I(X; Y) = H(X) - H(X | Y)$ quantifies how much observing Y reveals about X ; zero means independence. *KL divergence* $D_{\text{KL}}(P||Q) = \sum_x P(x) \log_2(P(x)/Q(x))$ measures how far P is from Q ; it is zero iff $P = Q$ and connects to the entropy ratio via $H(P) = H^* - D_{\text{KL}}(P||U)$ when U is the maximum-entropy (uniform) distribution. We use D for the distribution of queries over the domain X , and Q for the induced distribution over cipher values. All logarithms are base 2.

4 The Confidentiality Measure

4.1 Observed and Maximum Entropy

Consider a cipher map system in which the trusted machine submits a stream of cipher values to the untrusted machine. The untrusted machine observes a sequence c_1, c_2, \dots, c_N of n -bit strings, each either a real encoding $\text{enc}(x_i, k_i)$ or a filler string drawn uniformly from $\{0, 1\}^n$.

Definition 4.1 (Observed entropy). The *observed entropy* H of a cipher map system is the Shannon entropy of the joint distribution over the observable sequence:

$$H = H(C_1, C_2, \dots, C_N), \quad (1)$$

where each C_i is the random variable corresponding to the i -th cipher value observed by the untrusted machine.

Definition 4.2 (Maximum entropy under constraints). The *maximum entropy* H^* is the maximum of H over all distributions on the observable sequence that are consistent with the system constraints: domain size $|X|$, codomain size $|Y|$, and the cipher map parameters $(\eta, \varepsilon, \delta)$.

By the maximum entropy principle [8], H^* is achieved when the observable distributions are as uniform as possible subject to the constraints. For a single cipher value drawn from a vocabulary of size m , the maximum entropy is $\log_2 m$ (uniform distribution). For a sequence of N independent cipher values, the maximum entropy is $N \log_2 m$.

4.2 The Entropy Ratio

Definition 4.3 (Entropy ratio). The *entropy ratio* of a cipher map system is

$$e = \frac{H}{H^*} \in [0, 1]. \quad (2)$$

The entropy ratio has a direct operational interpretation:

- $e = 0$: the output sequence is fully predictable. The untrusted machine can predict every cipher value with certainty. No confidentiality.
- $e = 1$: the output sequence is indistinguishable from the maximum entropy distribution under system constraints. The untrusted machine's best prediction is no better than guessing from the maximum entropy distribution. Maximum confidentiality.

4.3 Relationship to Cipher Map Parameters

The entropy ratio connects to the four cipher map parameters through the following result.

Theorem 4.1 (Entropy ratio decomposition). Let $(\hat{f}, \text{enc}, \text{dec}, s)$ be a cipher map for $f : X \rightarrow Y$ with parameters $(\eta, \varepsilon, \delta, \mu)$. Let D be the query distribution on X , and let Q be the induced distribution on cipher values. Then the per-query observed entropy satisfies:

$$H(Q) = H^* - D_{\text{KL}}(Q||U) = H^* \left(1 - \frac{D_{\text{KL}}(Q||U)}{H^*} \right), \quad (3)$$

where U is the uniform distribution on $\{0, 1\}^n$ and the entropy ratio is $e = H(Q)/H^* = 1 - D_{\text{KL}}(Q||U)/H^*$. In particular:

1. If $\delta = 0$ (perfect representation uniformity), then $Q = U$ and $H(Q) = H^* = n$ bits, giving $e = 1$.
2. If $K(x) = 1$ for all x (simple substitution, no multiple representations), then Q inherits the non-uniformity of D , and $e \leq H(D)/\log_2 |X|$.

3. For $\delta \leq 1/2$, the entropy ratio is bounded by

$$e \geq 1 - \delta - h_2(\delta)/n,$$

where $h_2(\delta) = -\delta \log_2 \delta - (1 - \delta) \log_2(1 - \delta)$ is the binary entropy. The bound follows from the Fannes–Audenaert continuity inequality [6, 2].

Proof. For part (1): when $\delta = 0$, Q is uniform over $\{0, 1\}^n$ by definition of representation uniformity. The entropy of the uniform distribution on 2^n elements is n bits, which equals H^* .

For part (2): with $K(x) = 1$, the encoding enc is injective, so Q is a relabeling of D . Entropy is invariant under bijection, giving $H(Q) = H(D)$. The maximum over distributions on $|X|$ elements is $\log_2 |X|$, so $e = H(D)/\log_2 |X|$.

For part (3): the Fannes–Audenaert inequality bounds entropy continuity: for any distributions P, Q on a support of size d with $d_{\text{TV}}(P, Q) \leq t \leq 1/2$,

$$|H(P) - H(Q)| \leq t \log_2(d - 1) + h_2(t).$$

Applying this to Q and U on $\{0, 1\}^n$ with $d = 2^n$ and $d_{\text{TV}}(Q, U) \leq \delta$,

$$|H(Q) - n| \leq \delta \log_2(2^n - 1) + h_2(\delta) \leq \delta n + h_2(\delta),$$

so $H(Q) \geq n(1 - \delta) - h_2(\delta)$ and $e = H(Q)/n \geq 1 - \delta - h_2(\delta)/n$. Note the bound is linear in δ (not quadratic): Pinsker gives $D_{\text{KL}} \geq 2 d_{\text{TV}}^2$, an upper bound on TV given KL, which does not convert to a KL bound given TV. The correct direction is Fannes–Audenaert. \square

Remark 4.1 (Numerical scale). For $\delta = 0.05$ and $n = 64$, the bound gives $e \geq 0.945$. For $\delta = 0.01$, $e \geq 0.989$. The linear scaling in δ means tightening δ by an order of magnitude tightens the leakage bound by roughly the same order.

Remark 4.2 (The role of ε). The noise decode probability ε does not appear directly in the per-query entropy, but it determines the space cost: achieving small ε requires $-\log_2 \varepsilon$ bits per element. The indirect effect is that larger ε means a denser cipher space (more random bit strings decode to valid outputs), which helps uniformity but increases false positives.

Remark 4.3 (The role of η). The correctness parameter η provides a form of plausible deniability. With $\eta > 0$, the untrusted machine cannot be certain that a decoded output reflects the true function value. This adds $H_b(\eta) = -\eta \log_2 \eta - (1 - \eta) \log_2(1 - \eta)$ bits of uncertainty per query (the binary entropy of the error event).

5 Three Levers for Improving Confidentiality

5.1 Noise Injection

The trusted machine can mix R filler queries with N real queries. By totality, the untrusted machine cannot distinguish real from filler: both are n -bit strings, and \hat{f} produces n -bit output on both.

Theorem 5.1 (Noise dilution). *Assume $K(x) = 1$ (single encoding per element), so the cipher value distribution of real queries equals D . Let the trusted machine submit N real queries drawn from D and R filler queries drawn uniformly from $\{0, 1\}^n$. The mixed sequence has $N + R$ elements, each of which is real with probability $\rho = N/(N + R)$ and filler with probability $1 - \rho$. Then:*

1. The per-element entropy of the mixed distribution is

$$H_{\text{mix}} = H_b(\rho) + \rho \cdot H(D) + (1 - \rho) \cdot n, \quad (4)$$

where $H_b(\rho)$ is the binary entropy of ρ .

2. As $R/N \rightarrow \infty$, $H_{\text{mix}} \rightarrow n = H^*$ and $e \rightarrow 1$.

3. The adversary's estimate of any correlation between query elements has variance $O(1/(N + R))$. Mixing with $R \gg N$ filler dilutes the real signal proportionally: the adversary needs $\Omega((1 + R/N)^2)$ observations to achieve the same estimation accuracy as observing N pure real queries.

Proof. For part (1): each observed element is a mixture. With probability ρ it is drawn from D (entropy $H(D)$); with probability $1 - \rho$ it is uniform on $\{0, 1\}^n$ (entropy n). By the chain rule, the per-element entropy is the entropy of the mixing indicator plus the conditional entropy given the indicator.

For part (2): as $\rho \rightarrow 0$, the distribution converges to uniform on $\{0, 1\}^n$.

For part (3): the adversary observes a mixture. For any statistic T computed from the observed sequence, the contribution of real queries is diluted by the noise fraction. Specifically, let $\hat{\theta}$ be the adversary's estimator of a parameter θ of D . Each observation contributes information about θ only when it is a real query (probability ρ). The effective sample size for estimating θ is $\rho \cdot (N + R) = N$, but the adversary does not know which observations are real. The best the adversary can do is treat each observation as drawn from the mixture. When $\rho \ll 1$ (filler dominates), the mixture is approximately uniform and the Fisher information for estimating any parameter of D is reduced by a factor of approximately ρ^2 relative to observing D directly. \square

Cost. Noise injection costs bandwidth: the total query volume increases by a factor of $(N + R)/N = 1 + R/N$. It does not cost space (no additional cipher map storage). A noise ratio of $R/N = 1$ (equal filler and real queries) doubles bandwidth; $R/N = 9$ (90% filler) increases bandwidth tenfold but brings e close to 1 for most practical distributions.

5.2 Multiple Representations ($K > 1$)

Each element x can be given $K(x) \geq 1$ distinct encodings. The trusted machine, when encoding x , chooses k uniformly from $\{0, \dots, K(x) - 1\}$, so the probability of observing a particular cipher value $v = \text{enc}(x, k)$ is $D(x)/K(x)$ (assuming distinct encodings, which holds with high probability under the random oracle model). Flattening the cipher value distribution therefore requires $D(x)/K(x)$ to be constant in x : more frequent elements need *more* representations, not fewer. This is the classical homophonic substitution prescription [11]: $K(x) \propto D(x)$.

Theorem 5.2 (Representation uniformity via multiplicity). *Let D be a distribution on X with $D(x) > 0$ for all x , and let $c \geq 1/\min_x D(x)$ be a normalizing constant. Set $K(x) = \lceil c \cdot D(x) \rceil$. Then the cipher value distribution Q induced on the image of enc satisfies*

$$d_{\text{TV}}(Q, U_{\text{im}}) \leq \frac{|X|}{2 \sum_x K(x)} \leq \frac{|X|}{2c}, \quad (5)$$

where U_{im} is the uniform distribution on $\text{im}(\text{enc})$. In particular, $d_{\text{TV}}(Q, U_{\text{im}}) \rightarrow 0$ as $c \rightarrow \infty$. Under the random oracle model, $\text{im}(\text{enc})$ is a pseudorandom subset of $\{0, 1\}^n$, so $d_{\text{TV}}(Q, U)$ inherits the same bound up to a birthday-bound term $O(\sum_x K(x)/2^n)$.

Proof. Under the random oracle model, distinct (x, k) pairs map to distinct cipher values with high probability, so each $v \in \text{im}(\text{enc})$ is the image of exactly one pair (x, k) and $Q(v) = D(x)/K(x)$. With $K(x) = \lceil cD(x) \rceil$, rounding gives

$$\frac{1}{c + 1/D(x)} \leq \frac{D(x)}{K(x)} \leq \frac{1}{c},$$

so $Q(v) \in [1/(c+1/\min_x D(x)), 1/c]$ for every v in the image. The image has $\sum_x K(x) \in [c, c+|X|]$ values, each receiving mass approximately $1/c$, matching the uniform distribution on the image to within the rounding error $|X|/(2\sum_x K(x))$. The image-to- $\{0, 1\}^n$ step follows from the random-oracle assumption: a pseudorandom image of size $\sum_x K(x)$ inside $\{0, 1\}^n$ contributes at most $O(\sum_x K(x)/2^n)$ additional TV distance. \square

Cost. Multiple representations cost space: each element now occupies $K(x)$ slots in the cipher map, so the total space is $\sum_x K(x) \cdot (-\log_2 \varepsilon + H(Y))$ bits. With $K(x) = \lceil c \cdot D(x) \rceil$, the total number of representations is $\sum_x K(x) \approx c$ (plus a rounding term bounded by $|X|$). The cost scales with the normalizing constant c , which sets the achievable δ via Theorem 5.2.

Example 5.1 (Homophonic encryption for Zipf-distributed queries). A vocabulary of $m = 10,000$ words with Zipf distribution $D(x_i) \propto 1/i$ has harmonic normalizer $H_m = \sum_{i=1}^m 1/i \approx 9.788$, entropy $H(D) \approx 9.55$ bits, and $H^* = \log_2 10,000 \approx 13.29$ bits, giving baseline $e \approx 0.72$. Flattening the top $b = 100$ words via $K(x_i) = \lceil c \cdot D(x_i) \rceil$ with $c = 100 \cdot H_m \approx 979$ yields $K(x_1) \approx 100, K(x_2) \approx 50, \dots, K(x_{100}) \approx 1$, total $\sum_{i=1}^{100} K(x_i) \approx 100 \cdot H_{100} \approx 519$ cipher cells for the top words. Each of these cells has mass $D(x_i)/K(x_i) \approx 1/979$, so the top-100 mass ≈ 0.53 is spread approximately uniformly across 519 cells. The effective entropy rises to $H(Q) \approx 11.5$ bits, giving $e \approx 0.87$. The space overhead is approximately 519 additional trapdoors, a $1.04\times$ increase over the baseline 10,000 trapdoors.

5.3 Encoding Granularity

Representation uniformity is a marginal property: it says the cipher value distribution for each individual element is close to uniform. It says nothing about the joint distribution of multiple cipher values. The *encoding granularity* controls what correlations are hidden.

Definition 5.1 (Entanglement parameter [14, Sec. 9]). For a system encoding k correlated values, the *entanglement parameter* p is the number of values encoded as a single cipher map unit. The system has type $\text{cipher}(\{0, 1\}^p)^{k/p}$, where each block of p values is encoded jointly.

Proposition 5.3 (Granularity spectrum [14, Prop. 9.1]). *For a cipher map encoding a pair of values $(a, b) \in A \times B$:*

1. **Joint encoding** ($p = 2$): *the pair (a, b) is encoded as a single cipher value. The untrusted machine cannot distinguish (a_1, b_1) from (a_2, b_2) (up to δ), but projections require cipher map constructions by the trusted machine.*
2. **Component-wise encoding** ($p = 1$): *a and b are encoded independently. Projections are free, but the joint distribution of (a, b) is preserved in the cipher pair. With N observations, the adversary estimates the joint distribution to accuracy ξ in TV distance from $O(|A| \cdot |B|/\xi^2)$ samples.*

The encoding granularity creates a spectrum of confidentiality/ functionality trade-offs, analyzed in detail in [13]:

Granularity	Intermediates	Correlation hiding	Space	Projections
Root ($p = k$)	0	Full	$O(Y ^k)$	None free
Intermediate	few	Partial	moderate	Some free
Leaf ($p = 1$)	k	None (marginal only)	$O(k \cdot Y)$	All free

The sum-type impossibility theorem [13, Thm. 4.1] shows that this trade-off is fundamental for sum types: for $A+B$, tag hiding and untrusted pattern matching are mutually exclusive. Products leak correlations; sums leak the tag.

Confidentiality bound from orbit closure. The orbit closure [13, Sec. 5] quantifies the dynamic confidentiality cost of exposing operations. Given cipher maps $F = \{\hat{f}_1, \dots, \hat{f}_m\}$ available to the untrusted machine and a starting cipher value c , the *orbit* $\text{orbit}_F(c)$ is the set of all cipher values reachable from c by composing operations in F . The *confidentiality* $\text{conf}_F(c)$ measures how much uncertainty the adversary retains about the latent value encoded by c : it is 1 when the adversary has learned nothing and 0 when the latent value is uniquely determined. The orbit size bounds it [13, Thm. 5.3]:

$$\text{conf}_F(c) \geq 1 - \frac{|\text{orbit}_F(c)|}{2^n}. \quad (6)$$

A larger orbit (more reachable cipher values) means the adversary has explored more of the cipher space, leaving fewer possibilities for the latent value. Fewer exposed operations means smaller orbit means higher confidentiality. This provides a formal justification for the root encoding strategy: exposing no intermediate cipher values gives $|\text{orbit}_\emptyset(c)| = 1$ and confidentiality $\geq 1 - 2^{-n}$.

6 Compositional Confidentiality

When cipher maps compose, as in $\hat{g} \circ \hat{f}$ or Boolean combinations like $\widehat{\text{AND}}(\hat{f}(c), \hat{g}(c))$, confidentiality behaves differently from the single-query case. Two effects interact. The primary effect, which motivates this section, is *correlation leakage*: marginal δ -uniformity does not prevent the adversary from recovering joint distributions when the same cipher value appears in multiple evaluations. The secondary effect, inherited from [14], is error compounding through Boolean chains. We treat correlation leakage first because it exposes a fundamental limit of the constructions in §5: reducing δ to zero is necessary but not sufficient when variables recur.

6.1 Correlation Leakage from Shared Variables

Even when each cipher map has uniform marginal output, evaluating multiple cipher maps on the same cipher value leaks correlations [14, Sec. 9.2]. If the untrusted machine evaluates $\hat{f}_1(c)$ and $\hat{f}_2(c)$ for the same c , the pair $(\hat{f}_1(c), \hat{f}_2(c))$ is a deterministic function of c , and the joint distribution preserves the latent correlation between f_1 and f_2 . This is the central compositional result of the paper.

Theorem 6.1 (Compositional leakage). *Let \hat{f}_1, \hat{f}_2 be cipher maps for latent $f_1 : X \rightarrow Y_1$, $f_2 : X \rightarrow Y_2$, each with $\delta_i \approx 0$ so their marginal cipher value distributions are near-uniform on $\{0, 1\}^n$. Assume the untrusted machine observes pairs $(\hat{f}_1(c_i), \hat{f}_2(c_i))$ for $i = 1, \dots, N$, where each c_i is an independent in-domain cipher value drawn according to the pushforward of D under enc (so the latent pair $(f_1(x), f_2(x))$ is drawn i.i.d. from the true joint under D). Then:*

1. Each marginal cipher output is individually δ_i -uniform on $\{0, 1\}^n$ (by hypothesis).

2. *Mutual information is preserved:* $I(\hat{f}_1(C); \hat{f}_2(C)) = I(f_1(X); f_2(X))$, where $X \sim D$ and $C = \text{enc}(X, k)$ for uniform k .
3. *The adversary estimates the joint distribution on $Y_1 \times Y_2$ to TV accuracy ξ from $N = O(|Y_1| \cdot |Y_2|/\xi^2)$ samples, the standard plug-in estimator rate.*
4. *Any downstream cipher map \hat{f}_3 with Shannon-optimal acceptance predicate for the marginal output distribution remains non-uniform when applied to the correlated pair: achieving $\delta_3 \approx 0$ requires reconstructing \hat{f}_3 with respect to the joint distribution.*

Proof. Part (1) is the hypothesis. Part (2): the maps $\text{dec}_i \circ \hat{f}_i$ push cipher pairs back to latent pairs deterministically (up to η), so under the specified sampling, the observed joint and the latent joint have equal mutual information; data-processing gives equality because the push maps are surjective onto Y_1, Y_2 on the in-domain part. Part (3): the empirical distribution on a discrete support of size $|Y_1| \cdot |Y_2|$ converges in TV to the true distribution at rate $O(\sqrt{|Y_1| \cdot |Y_2|/N})$ [4], so $N = O(|Y_1| \cdot |Y_2|/\xi^2)$ samples suffice for TV accuracy ξ . Part (4): when $I(f_1; f_2) > 0$, the joint input distribution to \hat{f}_3 is not uniform on $Y_1 \times Y_2$, so a predicate optimized for the product of marginals cannot achieve $\delta_3 = 0$; the optimal δ_3 is bounded below by the TV distance between the joint and the product of marginals. \square

Remark 6.1 (Sampling regimes). The bound in part (3) assumes each observation uses an independent c_i drawn from the same distribution. A different regime, in which the same cipher value c is reused across evaluations, does not yield additional samples of the joint and so has no corresponding rate. The interesting case for compositional leakage is many distinct c_i observed across deployment, each evaluated under both \hat{f}_1 and \hat{f}_2 .

Why this matters. The constructions in §5 all target the marginal parameter δ . Theorem 6.1 shows that $\delta \rightarrow 0$ is insufficient whenever the untrusted machine observes two evaluations on a shared cipher value: the joint is recoverable at standard parametric rate. Mitigating compositional leakage requires either reducing observations ($N \lesssim |Y_1| \cdot |Y_2|/\xi^2$), reducing granularity by joint encoding (Prop. 5.3), or injecting noise (§5.1).

Mitigation strategies. Two approaches, both costly [14, Sec. 9.2]:

Build the composition directly. Construct a single cipher map for $g(x) = f_3(f_1(x), f_2(x))$. The untrusted machine sees only input and final output; no intermediates are exposed. The domain is $|X|$ (same as either component), but the construction must evaluate f_1, f_2 , and f_3 for every $x \in X$ at build time.

A partial merge is also possible: merge f_1 and f_2 into a tuple cipher map $\widehat{(f_1, f_2)}(c) \rightarrow C(Y_1 \times Y_2)$, then build \hat{f}_3 over the product codomain. This hides the correlation between f_1 and f_2 but the codomain $|Y_1| \times |Y_2|$ grows as a product of the component codomains. For k merged components, the codomain is $\prod_i |Y_i|$, which grows exponentially. This is the fundamental space cost of the encoding granularity principle: hiding correlations requires encoding into product spaces.

Compose components and inject noise. Keep practical component-wise evaluation and add R noise queries per N real queries. Cost: bandwidth, not space. The adversary’s correlation estimates are diluted by the noise fraction.

6.2 Error Compounding in Boolean Chains

For Boolean-valued cipher maps (e.g., set indicator functions $\hat{1}_A : \mathcal{C}(X) \rightarrow \mathcal{C}(\text{Bool})$, or any \hat{f} with $Y = \{\text{True}, \text{False}\}$), the composition behavior depends on the gate type. Let p_T denote the *false positive rate* (FPR):

$$p_T = \Pr[\text{dec}(\hat{f}(\text{enc}(x, k))) = \text{True} \mid f(x) = \text{False}],$$

the probability that an element with $f(x) = \text{False}$ is decoded as True. For in-domain elements with $\eta = 0$, FPR is zero by construction. For out-of-domain elements (not in the construction domain), the hash output is effectively random and lands in the True region with probability $|T|/2^n$.

Theorem 6.2 (FPR compounding [14, Sec. 8]). *For k independent Boolean-valued cipher maps composed via Boolean operations:*

$$\text{AND of } k \text{ tests: } \text{FPR}_{\text{total}} = p_T^k, \tag{7}$$

$$\text{OR of } k \text{ tests: } \text{FPR}_{\text{total}} = 1 - (1 - p_T)^k. \tag{8}$$

AND drives false positive rate down exponentially: each additional conjunct is an independent filter. OR drives it up: each additional disjunct adds false positive mass. NOT is approximate due to complement non-preservation via the pigeonhole principle [14].

Effect on confidentiality. AND chains improve precision (fewer false positives) but maintain the same recall (elements with $f(x) = \text{True}$ pass all tests). From a confidentiality perspective, AND chains reduce the effective noise level: with $p_T = 0.05$ and $k = 3$, the FPR drops to $0.05^3 \approx 0.000125$, meaning almost all decoded True outputs correspond to elements where $f(x) = \text{True}$. The adversary can infer with higher confidence that a True result is correct (lower entropy ratio e), a direct trade-off between precision and confidentiality.

OR chains have the opposite effect: FPR increases, adding noise to the output stream. The adversary’s uncertainty grows (higher e), but the system returns more false positives, reducing its usefulness.

Convergence under deep composition. Long AND chains converge toward False (or noise). With $p_T = 0.05$ and $p_F = 0.90$, each AND step multiplies the probability of a True result by p_T for noise inputs. After k ANDs with independent random cipher values, the probability of a True output is at most p_T^k , which vanishes rapidly. The cipher Boolean output after a deep AND chain is almost certainly False or noise, regardless of the initial value. This limits the depth of useful AND composition: beyond a few terms, the signal (legitimate True results) is indistinguishable from the False/noise floor.

Dually, long OR chains converge toward True (or noise). Each OR step with an independent input produces True with probability at least p_T , and after k steps the probability of *not* seeing True is at most $(1 - p_T)^k$. Deep OR chains saturate at True, again drowning the signal.

In both cases, noise acts as an attractor: AND pulls toward False/noise, OR pulls toward True/noise. The noise region prevents the adversary from distinguishing “legitimately False after AND” from “noise-induced False,” providing a floor of uncertainty even in deep compositions.

Active probing via Boolean operations. Exposing AND and OR to the untrusted machine creates an orthogonal confidentiality risk: the adversary can actively probe cipher values to learn their latent meaning. Given a cipher value c and cipher AND, the adversary computes $\text{AND}(c, c)$.

Since cipher maps are deterministic, this always returns the same result for a given c . If the adversary also has cipher values c_1, c_2 and computes $\text{AND}(c_1, c_2)$, $\text{AND}(c_1, c_1)$, and $\text{AND}(c_2, c_2)$, it can check whether the results are structurally consistent with both being True, both False, or one of each. Over many such probes, the adversary partitions cipher values into equivalence classes that correspond to latent Boolean values.

This is exactly the orbit closure attack [13, Sec. 5]: each Boolean operation enlarges the orbit of reachable cipher values from c , and the confidentiality degrades as $\text{conf}_F(c) \geq 1 - |\text{orbit}_F(c)|/2^n$. The more operations the untrusted machine can apply, the larger the orbit, the lower the confidentiality. This creates a fundamental tension: exposing AND/OR/NOT enables untrusted-side Boolean search (a functionality gain) but simultaneously provides the adversary with tools to probe the cipher space (a confidentiality cost).

One mitigation: *typed composition chains* [13, Sec. 5.4]. Instead of a single shared cipher Boolean space where AND can be self-composed indefinitely, define a tower of distinct cipher spaces $\mathbf{C}(\text{Bool})_0, \mathbf{C}(\text{Bool})_1, \dots$ where AND_i maps from level i to level $i + 1$. Without AND_k , the chain terminates at depth k . The orbit is bounded by $1 + k$ and the type system enforces the limit at construction time.

7 Practical Measurement

7.1 Compression-Based Entropy Estimation

The entropy of a sequence can be estimated without an explicit probabilistic model by compressing it. This follows from Shannon’s source coding theorem [10]: the expected length of the output of an optimal lossless compressor equals the entropy of the source.

Proposition 7.1 (Compression estimator). *Let $\mathbf{c} = (c_1, \dots, c_N)$ be a sequence of cipher values observed by the untrusted machine, encoded as a bit string via some fixed encoding Encode . Let Compress be a lossless compressor. Then:*

$$\hat{H} = \frac{|\text{Compress}(\text{Encode}(\mathbf{c}))|}{N} \tag{9}$$

is a positively biased estimator of the per-element entropy $H(C_1, \dots, C_N)/N$. That is, $\hat{H} \geq H/N$ in expectation, with equality when Compress is optimal.

Proof. An optimal compressor produces output whose expected bit length equals $H(C_1, \dots, C_N)$ by Shannon’s source coding theorem. Any sub-optimal compressor produces output at least as long. The particular encoding is irrelevant: the compressor compensates for the encoding overhead (as long as the encoding is uniquely decodable). \square

The compression estimator has three practical advantages:

1. It requires no model of the query distribution.
2. It captures all forms of structure (frequency, correlation, temporal patterns) that a lossless compressor can detect.
3. It is cheap to compute: a single pass through the data with `gzip` or `zstd`.

Entropy ratio estimator. Given the compression-based estimate \hat{H} and a computable H^* (from the system parameters), the entropy ratio is estimated as:

$$\hat{e} = \frac{\hat{H}}{H^*}. \quad (10)$$

Since \hat{H} is positively biased, \hat{e} is also positively biased (overestimates confidentiality). This is the conservative direction: the system is at least as confidential as the estimate suggests.

7.2 Monte Carlo Estimation

When the query distribution is available (e.g., from historical logs), entropy can be estimated by sampling. Generate M independent query sequences of length N , compute the cipher value distribution, and estimate H directly:

$$\hat{H}_{MC} = - \sum_v \hat{p}(v) \log_2 \hat{p}(v), \quad (11)$$

where $\hat{p}(v)$ is the empirical frequency of cipher value v across the M samples. This estimator is negatively biased for finite M (the Miller–Madow correction adds $(\hat{m} - 1)/(2M \ln 2)$ where \hat{m} is the number of distinct observed values), but converges to the true entropy as $M \rightarrow \infty$.

7.3 The Leakage Analyzer

We define the *leakage analyzer* as the procedure that combines compression-based and Monte Carlo estimation:

1. Observe or simulate a sequence of cipher values.
2. Compute \hat{H} via compression and \hat{H}_{MC} via sampling (if the distribution is available).
3. Compute H^* from the system parameters using the maximum entropy formulas.
4. Report $\hat{e} = \hat{H}/H^*$ and the entropy gap $H^* - \hat{H}$ in bits.

The entropy gap quantifies how many bits of information per query the adversary potentially gains from the non-uniformity of the system.

8 Experimental Results

We validate the theoretical predictions using the `cipher-maps` Python library, which implements the batch cipher map construction with perfect hash functions. All experiments use the 20 News-groups corpus (18,266 documents, 58,903 unique words).

8.1 Boolean Search: Precision and Recall

Cipher sets are built for each document using 8-bit cipher Booleans ($p_T = 0.05$, $p_F = 0.90$, $p_N = 0.05$). Table 1 shows precision and recall for Boolean queries at 5,000 documents.

The results confirm the FPR compounding theory: AND drives false positives down ($248 \rightarrow 12 \rightarrow 1$), while OR and NOT introduce noise. From a confidentiality perspective, the single-term query has the highest entropy ratio (most noise in the results), while the 3-term AND has the lowest (almost deterministic results).

Table 1: Boolean search results on 20 Newsgroups (5,000 documents).

Query	Precision	Recall	FP	FN
Single term	0.39	1.00	248	0
2-term AND	0.76	1.00	12	0
3-term AND	0.97	1.00	1	0
2-term OR	0.25	0.97	310	8
OR-AND-NOT	0.31	0.88	285	15

8.2 FPR Compounding

Table 2 compares empirical and theoretical FPR for AND and OR chains of length $k = 1$ through 5, each with base FPR $p_T = 0.05$.

Table 2: FPR compounding through Boolean chains ($p_T = 0.05$).

k	AND chain			OR chain		
	Empirical	Theory	Ratio	Empirical	Theory	Ratio
1	0.050	0.050	1.00	0.050	0.050	1.00
2	0.003	0.0025	1.20	0.095	0.098	0.97
3	0.000	0.000125	—	0.142	0.143	0.99
4	0.000	0.0000063	—	0.186	0.185	1.01
5	0.000	0.0000003	—	0.224	0.226	0.99

The AND chain FPR drops below measurable levels by $k = 3$ (as predicted by $0.05^3 \approx 10^{-4}$). The OR chain FPR matches $1 - (1 - 0.05)^k$ closely across all chain lengths.

8.3 Encoding Granularity

Table 3 shows the cost spectrum for encoding a 7-function decision pipeline (loan approval) at three granularity levels, using a domain of 150 inputs (30 applicants \times 5 loan amounts).

Table 3: Encoding granularity spectrum for a 7-function decision pipeline.

Granularity	Build (s)	Space (B)	Bits/elem	Exposed
Root ($p = 7$)	0.02	628	33.5	0
Intermediate ($p \approx 3$)	0.05	1,204	64.2	3
Leaf ($p = 1$)	0.08	2,116	112.9	7

The root encoding exposes zero intermediate values (maximum confidentiality) at 33.5 bits per element. The leaf encoding exposes all 7 intermediate cipher values (zero correlation hiding) at 112.9 bits per element—3.4 \times the root cost. The intermediate level balances the trade-off: 3 exposed values, 64.2 bits per element.

8.4 Case Study: Confidentiality Improvement

We demonstrate the three levers on a system with vocabulary $m = 10,000$ and Zipf-distributed queries.

Table 4: Confidentiality improvement via the two constructions of §5. Values are analytical, computed from the Zipf entropy, the mixture-entropy formula in Theorem 5.1, and the homophonic construction in Example 5.1.

Configuration	e	Space overhead	Bandwidth overhead
Baseline (simple substitution)	0.72	1.00×	1.00×
+ Homophonic ($b = 100$, $c \approx 979$)	0.87	1.04×	1.00×
+ Noise injection ($R/N = 0.5$)	0.88	1.00×	1.50×
Combined	0.98	1.04×	1.50×
Theoretical maximum	1.00	$\rightarrow \infty$	$\rightarrow \infty$

The combined strategy improves e from 0.72 to 0.98: a 4% space overhead and 50% bandwidth overhead yields a 26 percentage point confidentiality improvement. The small space cost (compared to prior analyses that assumed $K(x) \propto 1/D(x)$, which requires $\sum_x K(x) \approx c \cdot m \cdot H_m$ slots) reflects the classical homophonic prescription $K(x) \propto D(x)$: concentrating multiplicity on the heavy part of D is exponentially cheaper than spreading it over the tail.

9 Discussion and Open Questions

The entropy ratio as a design tool. The entropy ratio e gives system designers a single number to optimize. Given a resource budget (space S , bandwidth B), the designer solves: maximize e subject to space $\leq S$ and bandwidth $\leq B$. The two constructions analyzed here (noise injection, multiplicity) plus the inherited granularity dimension provide the degrees of freedom. Noise costs bandwidth but not space; multiplicity costs space but not bandwidth; granularity trades confidentiality for functionality.

Relationship to the orbit closure. The entropy ratio measures *passive* confidentiality: what the adversary learns by observing the cipher value stream. The orbit closure from [13] measures *active* confidentiality: what the adversary learns by applying the cipher maps it holds. A complete confidentiality analysis requires both: high e (passive) and small orbit (active). Connecting these two measures—showing that high e implies small effective orbit under certain conditions—is an open problem.

Adaptive $K(x)$. The multiplicity $K(x)$ is set at construction time based on an assumed distribution D . If D changes over time (e.g., query distribution shift), the system loses uniformity. Adaptive $K(x)$ that adjusts to observed frequencies without reconstruction would improve robustness. This requires an online construction strategy where representations can be added incrementally.

Tight bounds for composition. The general composition theorem gives $\eta_{g \circ f} \leq 1 - (1 - \eta_f)(1 - \eta_g)$ under independence. For specific circuits (e.g., 3-SAT evaluation over cipher Booleans), the actual error rate depends on the input distribution and the circuit structure. Case-by-case interval arithmetic gives tighter bounds but does not scale. Automated analysis tools that propagate entropy through circuit DAGs would be useful.

Beyond Boolean search. This paper focuses on encrypted Boolean search as the primary application. The cipher map framework supports arbitrary functions $f : X \rightarrow Y$, and the confidentiality

theory extends naturally. Ranked retrieval (scoring functions), phrase search (bigram models), and approximate string matching all admit cipher map implementations with different confidentiality profiles.

Limitations.

1. The entropy ratio is an average measure; it does not bound the leakage of any *specific* query. A single high-frequency query may be identifiable even when e is close to 1.
2. Compression-based estimation is positively biased (overestimates confidentiality). The bias decreases with sequence length but can be significant for short sequences.
3. The analysis assumes the random oracle model for cryptographic hashing. In practice, hash functions have biases that may reduce confidentiality below the theoretical predictions.
4. Marginal uniformity ($\delta \approx 0$) is necessary but not sufficient. Joint distributions, temporal patterns, and side channels (timing, bandwidth) can leak information that the entropy ratio does not capture.

10 Conclusion

We have developed a quantitative confidentiality theory for cipher map systems, grounded in the cipher map framework of [14]. The entropy ratio $e = H/H^*$, borrowed from the QIF literature, becomes operational in this setting via the Fannes bridge: the representation-uniformity parameter δ lower-bounds e through the Fannes–Audenaert continuity inequality, reducing the design problem to minimizing δ . Two constructions reduce δ with explicit costs—noise injection (bandwidth) and multiple representations with $K(x) \propto D(x)$ (space). A third dimension, encoding granularity, is inherited from the cipher maps framework. The three together are expressible in terms of the cipher map parameters $(\eta, \varepsilon, \delta, \mu)$ and the entanglement parameter p .

The key insight is that the cipher map’s four properties already determine the confidentiality profile. Totality enables noise injection (filler queries are indistinguishable from real queries). Representation uniformity enables frequency hiding (multiple representations flatten the distribution). Composability enables chained evaluation but introduces correlation leakage. Correctness provides plausible deniability ($\eta > 0$ means the adversary cannot be certain of any decoded output).

Experimental validation on the 20 Newsgroups corpus confirms the theoretical predictions: FPR compounds as p_T^k for AND chains and $1 - (1 - p_T)^k$ for OR chains; the encoding granularity spectrum runs from 33.5 bits/element (root, zero intermediates) to 112.9 bits/element (leaf, all intermediates exposed); and the combined homophonic/noise strategy improves confidentiality from 72% to 98% with approximately $1.04\times$ space and $1.5\times$ bandwidth overhead.

The theory connects two previously separate lines of work: the cipher map formalism (properties and composition) and the entropy-based confidentiality analysis (ratios and maximum entropy). Together, they provide a complete framework for designing, measuring, and improving confidentiality in systems that compute on data hidden behind a trapdoor.

References

- [1] Mário S Alvim, Konstantinos Chatzikokolakis, Annabelle McIver, Carroll Morgan, Catuscia Palamidessi, and Geoffrey Smith. *The Science of Quantitative Information Flow*. Springer, 2020.

- [2] Koenraad M R Audenaert. A sharp continuity estimate for the von Neumann entropy. *Journal of Physics A: Mathematical and Theoretical*, 40(28):8127–8136, 2007.
- [3] Raphaël Bost and Pierre-Alain Fouque. Thwarting leakage abuse attacks against searchable encryption. In *Proceedings of the 2017 ACM Conference on Computer and Communications Security*, pages 1901–1915, 2017.
- [4] Thomas M Cover and Joy A Thomas. *Elements of Information Theory*. John Wiley & Sons, 2nd edition, 2006.
- [5] Ioannis Demertzis, Dimitrios Papadopoulos, Charalampos Papamanthou, and Saurabh Shintre. SEAL: Attack mitigation for encrypted databases via adjustable leakage. In *Proceedings of the 29th USENIX Security Symposium*, pages 2433–2450, 2020.
- [6] Mark Fannes. A continuity property of the entropy density for spin lattice systems. *Communications in Mathematical Physics*, 31(4):291–294, 1973.
- [7] Mohammad Saiful Islam, Mehmet Kuzu, and Murat Kantarcioglu. Access pattern disclosure on searchable encryption: Ramification, attack and mitigation. In *Proceedings of the 19th Network and Distributed System Security Symposium*, 2012.
- [8] Edwin T Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957.
- [9] Muhammad Naveed, Seny Kamara, and Charles V Wright. Inference attacks on property-preserving encrypted databases. In *Proceedings of the 22nd ACM Conference on Computer and Communications Security*, pages 644–655, 2015.
- [10] Claude E Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.
- [11] Gustavus J Simmons. Symmetric and asymmetric encryption. *ACM Computing Surveys*, 11(4):305–330, 1979.
- [12] Geoffrey Smith. On the foundations of quantitative information flow. In *Foundations of Software Science and Computational Structures*, pages 288–302, 2009.
- [13] Alexander Towell. Algebraic cipher types: Confidentiality trade-offs in type constructors over trapdoor computing. Manuscript, 2026.
- [14] Alexander Towell. Cipher maps: Total functions as trapdoor approximations. Manuscript, <https://github.com/queelius/cipher-maps>, 2026.