Beyond Shannon: Operational Perfect Secrecy as a Generalised Model for Information-Theoretic Security

Formalising Secrecy as Adversarial Success Probability

Adrian Neal
Oxford Scientifica

adrian.neal@oxfordscientifica.com

July 2025

Abstract

Shannon's 1949 theorem defines perfect secrecy as a condition where every possible message remains equally likely given any ciphertext, which requires a key at least as long as the message. This definition, while foundational, is binary and assumes uniform message priors—assumptions rarely met in real communication systems. It cannot express the fact that secrecy degrades gradually as key entropy decreases, and it does not account for semantic structure or contextual knowledge available to adversaries.

This paper extends Shannon's framework by introducing *Operational Perfect Secrecy (OPS)*, which defines secrecy in terms of adversarial success probability rather than requiring complete message-space coverage. Within this framework we also define two new forms of information-theoretic security: *Combinatorial ITS (C-ITS)*, which achieves OPS through combinatorial ambiguity of candidate decryptions, and *Dimensional Ambiguity ITS (DA-ITS)*, which achieves OPS by concealing the dimensionality of the key space itself. We show that OPS converges to Shannon secrecy when the support size approaches the message space, while providing meaningful guarantees even with shorter keys.

These results generalise the concept of perfect secrecy into a continuous, operational measure and establish a new theoretical foundation for scalable information-theoretic security.

Keywords: information-theoretic security; perfect secrecy; Shannon; operational security; adversarial success probability; combinatorial security; OPS; C-ITS; DA-ITS

1 Introduction

Information-theoretic security (ITS) represents the strongest known form of cryptographic protection. Unlike computational security, which assumes that adversaries are limited by current algorithmic knowledge and available computational power, ITS guarantees hold even

against an adversary with unbounded resources. Claude Shannon's 1949 framework formalised this notion, proving that a cipher achieves perfect secrecy if, and only if, its key is as long as the message, chosen uniformly at random, used only once, and kept perfectly secret [11]. This model is exemplified by the one-time pad (OTP), which remains the only widely acknowledged system proven to achieve perfect secrecy.

However, Shannon's model also imposes conditions that are rarely, if ever, satisfied in practice. It assumes that messages are drawn uniformly from the entire message space, while real messages are highly structured and semantically constrained. It assumes that keys are drawn from an idealised source of perfect randomness, while physical key sources are finite and imperfect. It treats secrecy as a binary property—either perfect or broken—which means security collapses completely if even a single condition is relaxed. These limitations are not practical inconveniences but conceptual: Shannon's definition cannot express that secrecy degrades gradually as key entropy decreases or as message structure becomes known.

This paper addresses this gap by generalising Shannon's framework. We propose *Operational Perfect Secrecy (OPS)*, a new information-theoretic definition that measures secrecy in terms of an adversary's optimal success probability rather than as a binary property of message posteriors. OPS subsumes Shannon secrecy as a special case while extending it to cover realistic conditions where message distributions are structured, keys are short, and randomness sources are imperfect. Within this framework we define two new families of ITS schemes, *Combinatorial ITS (C-ITS)* and *Dimensional Ambiguity ITS (DA-ITS)*, which achieve OPS using short keys and ephemeral public random blocks.

Contributions.

- We formalise the implicit assumptions underlying Shannon's secrecy theorem and show why they are incompatible with real communication systems.
- We introduce *Operational Perfect Secrecy (OPS)*, a new definition of ITS based on bounding adversarial success probability, and prove that it generalises Shannon secrecy.
- We define two new families of ITS—*Combinatorial ITS (C-ITS)* and *Dimensional Ambiguity ITS (DA-ITS)*—and prove that both achieve OPS under specified conditions.
- We compare OPS to existing secrecy notions such as ϵ-perfect secrecy and entropic security, showing how OPS resolves the discontinuity and impracticality of Shannon's model.

OPS reconceptualises secrecy as a graded property quantified by an adversary's success bound, rather than as a binary condition. This generalisation overcomes the discontinuity in Shannon's model and establishes a theoretical basis for constructing scalable and deployable information-theoretic security systems.

Relationship to follow-up work. This paper focuses exclusively on the theoretical framework of Operational Perfect Secrecy (OPS). A companion paper will present a concrete system, called *Q-Stream*, that implements OPS using large shared quantum-random blocks and short secret keys. That systems paper [9] builds on the definitions and proofs

developed here to show how OPS can be achieved in practical communication environments, but its design and performance aspects are beyond the scope of this work.

2 Background

2.1 Shannon Perfect Secrecy

Shannon defined *perfect secrecy* as the property that observing a ciphertext reveals no information about the underlying message. Formally, for message random variable M, key random variable K, and ciphertext random variable $C = E_K(M)$, a cipher has perfect secrecy if and only if

$$P(M = m \mid C = c) = P(M = m) \qquad \forall m, c.$$

This implies that the mutual information is zero:

$$I(M; C) = 0.$$

Shannon proved that perfect secrecy is achieved if, and only if, the following conditions hold:

- **Key length $\geq$ message length:** $|K| \geq |M|$
- **Uniform random key:** $P(K = k) = 1/|\mathcal{K}|$ for all k
- **Independence:** K is independent of M
- **One-time use:** each key is used exactly once
- **Reversibility:** given K and C, decryption uniquely recovers M

The one-time pad satisfies all of these conditions and is the canonical example of a perfectly secret cipher.

2.2 What Shannon Does Not Model

While Shannon's proof is mathematically sound, it relies on several implicit idealisations that do not hold in real communication systems:

- **Uniform message priors.** The proof assumes that messages are drawn uniformly from the full message space $\{0, 1\}^n$. Real messages are highly structured, predictable, and often known to the adversary to belong to a very small semantic subset.
- **Idealised random key source.** The proof models K only as an abstract uniform random variable. It does not discuss how this randomness is physically generated, implicitly assuming an oracle-like entropy source that can output perfectly uniform bits on demand.

- **No modelling of side information.** Contextual knowledge (e.g. protocol headers, natural language structure, metadata) is not considered. In reality, such side information lets an adversary discard most candidates even when the ciphertext is OTP-encrypted.
- **Binary security property.** The model defines secrecy as an all-or-nothing condition. If any of Shannon's assumptions are slightly violated—shorter keys, biased randomness, structured messages—the model provides no way to quantify residual secrecy.

These omissions mean Shannon secrecy is mathematically correct but not operationally usable as a design goal for modern systems. They motivate the need for a generalised secrecy model that tolerates structured messages, bounded-entropy keys, and imperfect randomness—while still providing provable information-theoretic guarantees.

3 Limitations of Shannon's Model

Although Shannon's secrecy theorem is mathematically sound, its applicability depends on several assumptions that do not hold in real communication systems. This section outlines four fundamental limitations that motivate the need for a more general secrecy framework.

3.1 Non-uniform Message Priors

Shannon's proof assumes that messages are drawn uniformly from the entire message space $\{0,1\}^n$, or at least that the message distribution is independent of the ciphertext. In practice, real messages are highly structured and exhibit significant redundancy (e.g. natural language, file formats, protocol fields). An adversary who knows this structure can eliminate most candidate decryptions as implausible, even if they have the correct ciphertext under a one-time pad. This violates the key assumption $P(M \mid C) = P(M)$ and means that Shannon secrecy cannot model the effect of semantic or structural knowledge.

3.2 Discontinuity of the Security Property

Shannon's definition treats secrecy as a binary property: it is either perfect or completely broken. This creates a discontinuity: if a system offers 2^n possible decryptions per ciphertext, it is perfectly secret, but if it offers only $2^n - 1$, it is not secret at all. This sharp threshold is conceptually implausible and does not reflect how security degrades in practice, where each reduction in key entropy or increase in adversary knowledge only gradually increases the probability of a correct guess. Shannon's model provides no way to quantify this partial security.

3.3 Contextual Leakage

The model assumes that the ciphertext is the adversary's only information about the message. In practice, adversaries often know metadata, expected formats, languages, timing, or communication context. This auxiliary knowledge collapses the effective candidate space, allowing the adversary to filter decryptions and identify plausible messages with much less

work than Shannon's analysis assumes. The classic "ambassadorial channel" scenario illustrates this: even if ciphertext is OTP-encrypted, the adversary can discard decryptions that are not well-formed diplomatic text, breaking the assumption of uniform message priors.

3.4 Randomness–Completeness Contradiction

Finally, Shannon's theorem implicitly conflates two incompatible requirements on the keyspace. It assumes both that:

(a) the keyspace contains all 2^n possible n-bit keys (completeness), and

(b) each key is "truly random," i.e. has full entropy or incompressible structure.

These conditions cannot both hold. A set containing all 2^n possible keys must also contain low-complexity keys (e.g. 0^n), which are not random and have zero entropy as fixed values. Conversely, a set restricted only to keys that are individually incompressible cannot be complete. This creates a logical inconsistency: Shannon's model requires a keyspace that cannot exist as defined. While this is hidden by treating K purely as a random variable, it means the model assumes an idealised entropy oracle rather than a realisable key source.

Taken together, these limitations mean that Shannon's perfect secrecy is valid only under idealised and unrealistic conditions. It cannot express partial secrecy or quantify adversarial success probability when keys are short, messages are structured, or auxiliary knowledge exists. This motivates the need for a generalised definition—*Operational Perfect Secrecy (OPS)*—which can model secrecy as a continuous property and apply to real systems.

4 Operational Perfect Secrecy (OPS)

The limitations described in Section 3 highlight that Shannon's notion of perfect secrecy, while mathematically correct, relies on idealised assumptions that do not hold in practice. It presumes uniform message priors, treats secrecy as a binary property, ignores contextual side information, and implicitly assumes an impossible keyspace that is both complete and entirely random. As a result, Shannon's model cannot describe how secrecy degrades when keys are shorter than messages, when messages are structured, or when adversaries have auxiliary knowledge.

To address these gaps, we introduce *Operational Perfect Secrecy (OPS)*, which reframes secrecy not as a property of posterior message distributions but as an explicit bound on the adversary's optimal success probability. OPS generalises Shannon's model, tolerates non-uniform message distributions, and provides a continuous measure of security even when the key length is much shorter than the message length.

Definition 1 (Operational Perfect Secrecy)**.** *An encryption scheme achieves* t-bit Operational Perfect Secrecy (OPS) *if for all adversaries* A*,*

$$\max_A \Pr[A(C) = M] \leq 2^{-t}.$$

Here t can be interpreted as the *operational secrecy level* in bits. Intuitively, it expresses that the ciphertext leaves at least 2^t messages indistinguishable from the adversary's perspective. This moves secrecy from an all-or-nothing property to a continuous quantity: if $t = |M|$, the scheme achieves perfect secrecy in the Shannon sense; if $t < |M|$, the adversary still faces 2^t plausible candidates and gains no useful advantage.

Theorem 1 (Support-size bound for OPS)**.** *If for every ciphertext c the set of messages m such that $\exists k\ E_k(m) = c$ has size at least 2^t, then the scheme achieves t-bit OPS.*

Proof sketch. Let $\mathcal{S}(c) = \{m : \exists k\ E_k(m) = c\}$. If $|\mathcal{S}(c)| \geq 2^t$, the optimal adversary can do no better than guessing uniformly among these candidates. The maximum success probability is thus at most $1/|\mathcal{S}(c)| \leq 2^{-t}$. □

Remark 1 (Relation to Shannon secrecy)**.** *If $t = |M|$ then $|\mathcal{S}(c)| = 2^{|M|}$ and each ciphertext is consistent with every possible message. This is exactly Shannon's perfect secrecy condition. Hence OPS strictly generalises Shannon secrecy: Shannon secrecy corresponds to the special case $t = |M|$.*

OPS therefore resolves the discontinuity of Shannon secrecy. Instead of collapsing from full security to zero when $|K|$ drops below $|M|$, security now degrades smoothly as t decreases. This makes OPS applicable to systems that use short keys, public randomness, or structured message spaces while still providing quantifiable information-theoretic security guarantees.

Remark 2 (How OPS resolves Shannon's limitations)**.** *OPS addresses the four limitations identified in Section 3 as follows:*

- ***Non-uniform message priors:*** *OPS does not assume any prior distribution on messages. It bounds the adversary's success probability directly, so it remains valid even when the message distribution is highly structured or known.*
- ***Discontinuity:*** *OPS makes secrecy a continuous quantity. Reducing key entropy or increasing adversarial knowledge gradually decreases the parameter t, rather than collapsing security from full to none as in Shannon's model.*
- ***Contextual leakage:*** *OPS explicitly accounts for all adversarial knowledge—context, metadata, language structure—by defining security solely in terms of the adversary's overall success probability given all information they hold.*
- ***Randomness–completeness contradiction:*** *OPS does not require the keyspace to contain all possible keys or that every key be individually "truly random." It only requires that enough plausible decryptions remain consistent with each ciphertext to keep the adversary's success probability below 2^{-t}.*

In this way, OPS generalises Shannon's secrecy condition from an idealised and discontinuous property into an operational measure that applies to real systems.

5 New ITS Families: C-ITS and DA-ITS

Operational Perfect Secrecy (OPS) provides a general target: an adversary's optimal success probability must be bounded by 2^{-t}. We now define two concrete families of schemes that satisfy OPS without requiring the key to be as long as the message. These constructions exploit ephemeral public random blocks together with short secret indices, and they are particularly relevant to practical deployments.

5.1 Combinatorial ITS (C-ITS)

C-ITS achieves OPS by ensuring that a short secret key can map into a combinatorially large number of possible pads extracted from an ephemeral public randomness block. The adversary can see the ciphertext and the randomness, but without the secret index they cannot determine which of the many possible pads was used.

Definition 2 (Combinatorial ITS (C-ITS))**.** *Let $Q \in \{0,1\}^N$ be a public randomness block and let $D \in \{0,1\}^d$ be a secret key. Let $F : \{0,1\}^d \times \{0,1\}^N \to \{0,1\}^n$ be an extraction function. Define the* extraction set

$$\mathcal{K}(Q) = \{F(d, Q) \mid d \in \{0,1\}^d\}$$

and the combinatorial richness

$$\kappa(Q) = \lfloor \log_2 |\mathcal{K}(Q)| \rfloor.$$

A scheme achieves k-bit C-ITS if $k = \min(d, \kappa(Q))$.

Theorem 2 (Security of C-ITS)**.** *If a scheme achieves k-bit C-ITS, then any adversary's success probability in recovering M from C is at most 2^{-k}.*

Proof sketch. Each ciphertext $C = M \oplus F(D, Q)$ is consistent with $|\mathcal{K}(Q)| \geq 2^k$ possible pads and thus 2^k possible messages. Without D, the adversary's optimal strategy is to guess among these candidates, giving success probability at most 2^{-k}. □

5.2 Dimensional Ambiguity ITS (DA-ITS)

DA-ITS extends this idea by also concealing the key's dimensionality — the adversary does not know how long the extracted pad is or which subspace it came from. This prevents the adversary from even constructing the correct probability space of candidates, further reducing their effective success probability.

Definition 3 (Dimensional Ambiguity ITS (DA-ITS))**.** *Let $\{\mathcal{D}_\ell\}$ be a family of possible key spaces of dimension ℓ, and let $\mathcal{K}_\ell(Q) = \{F(d, Q) \mid d \in \mathcal{D}_\ell\}$. Define*

$$k_\ell = \lfloor \log_2 |\mathcal{K}_\ell(Q)| \rfloor.$$

If the adversary does not know ℓ and for each admissible ℓ it holds that $|\mathcal{K}_\ell(Q)| \geq 2^{\min(\ell, k_\ell)}$, then the scheme achieves t-bit DA-ITS where

$$t = \min_\ell \min(\ell, k_\ell).$$

Theorem 3 (Security of DA-ITS)**.** *If a scheme achieves t-bit DA-ITS, then any adversary's success probability in recovering M from C is at most 2^{-t}.*

Proof sketch. Each ciphertext is consistent with at least $2^{\min(\ell,k_\ell)}$ candidates for every admissible ℓ. Since the adversary does not know which ℓ applies, they cannot normalise a posterior over the combined candidate set and can do no better than uniform guessing, giving success probability at most 2^{-t}. □

Remark 3. *While C-ITS relies purely on the number of candidate messages, DA-ITS adds a second layer of uncertainty by obscuring the dimensionality of the extraction space. This means the adversary must not only guess which candidate is correct but also which space it came from.*

Remark 4 (How C-ITS and DA-ITS instantiate OPS)**.** *The C-ITS and DA-ITS families provide concrete constructions that realise OPS using short keys and ephemeral public random blocks:*

- ***C-ITS:*** *Provides OPS through combinatorial ambiguity. A short key D selects and orders bits from a large public block Q, producing one of $|\mathcal{K}(Q)|$ possible pads. The adversary faces 2^k indistinguishable candidates, giving success probability $\leq 2^{-k}$ with $k = \min(d, \kappa(Q))$.*
- ***DA-ITS:*** *Extends C-ITS by also concealing the dimensionality ℓ of the key space. The adversary must guess not only which candidate is correct but which space it came from. This multiplies their uncertainty and yields a success bound $\leq 2^{-t}$ with $t = \min_\ell \min(\ell, k_\ell)$.*

Both families therefore satisfy OPS: they guarantee that each ciphertext leaves at least 2^t plausible candidates from the adversary's perspective, even when $d \ll n$ and the message distribution is fully known.

6 Comparison and Implications

The previous section presented two concrete families—C-ITS and DA-ITS—that achieve Operational Perfect Secrecy (OPS) using short keys and ephemeral public random blocks. These constructions demonstrate that OPS is not only a theoretical generalisation but also a practically realisable security property.

We now compare OPS with existing secrecy notions, including Shannon-perfect secrecy and entropic or ϵ-secrecy, to clarify how OPS fits within the broader landscape of information-theoretic security.

Operational Perfect Secrecy (OPS) can be seen as a generalisation of Shannon secrecy and as a complementary alternative to other relaxed information-theoretic notions such as ϵ-perfect secrecy and entropic security. Table 1 summarises the key distinctions.

OPS differs from prior notions in two central ways. First, it removes Shannon's implicit assumption of uniform message priors, allowing the security definition to remain valid even when messages are fully known to be structured or predictable. Second, it decouples secrecy

Property	**Shannon Perfect Secrecy**	**Entropic / ϵ-Secrecy**	**Operational Perfect Secrecy (OPS)**
Core Definition	$P(M \mid C) = P(M)$	$SD(P(M\|C), P(M)) \leq \epsilon$ or $\Pr[A(C) = f(M)] \approx \Pr[A(f(M))]$	$\max_A \Pr[A(C) = M] \leq 2^{-t}$
Assumptions	Key is uniform, secret, used once; $\|K\| \geq \|M\|$; messages uniform	Message has high min-entropy from adversary's view	Only requires many plausible candidates remain (support size $\geq 2^t$)
Secrecy Quantification	Binary (all-or-nothing)	Approximate (statistical distance or predicate advantage)	Continuous (success probability bound 2^{-t})
Key Length Requirement	$\geq \|M\|$	Can be shorter if message entropy is high	Tunable independently of $\|M\|$ via $t = \min(d, \kappa(Q))$
Message Prior Knowledge Allowed	None (assumes uniform)	Partial (bounded entropy loss tolerated)	Full (structured distributions tolerated)
Typical Use Case	OTP theoretical ideal	High-entropy data encryption with leakage bounds	Systems with large public randomness and short secret seeds
Status	Classical (1949)	Well-studied	New (this work)

Table 1: Comparison of secrecy notions. OPS generalises Shannon secrecy and tolerates structured messages and bounded key lengths.

from key length: rather than requiring the key entropy to meet or exceed message entropy, OPS expresses secrecy directly as a measurable bound on adversarial success probability. This reframes secrecy as a continuous property rather than a binary one, resolving the discontinuity at $|K| = |M|$ in Shannon's framework.

This shift has practical consequences. Because OPS allows short keys provided they create a large enough candidate set, it enables new system architectures that use ephemeral public random blocks (such as quantum-random blocks) with only small secret seeds. The C-ITS and DA-ITS families defined in Section 5 instantiate this principle. They achieve t-bit OPS even with $d \ll n$, making it possible to deliver information-theoretic security at

scale—something long thought impossible under Shannon's conditions.

Finally, while OPS generalises Shannon secrecy, it does not contradict it. When $t = |M|$, OPS collapses back to Shannon-perfect secrecy. OPS should therefore be viewed as a superset notion: it preserves Shannon's guarantees in the ideal case while extending them to realistic message distributions and finite-entropy key sources.

7 Related Work

Shannon's 1949 framework formalized perfect secrecy as $P(M \mid C) = P(M)$ and showed it is achieved by the one-time pad when the key is uniform, independent, used once, and at least as long as the message [11]. Shannon also analyzed message redundancy and the unicity distance, observing that real messages are highly non-uniform and that side-information can collapse candidate sets even under one-time pad encryption.

A large body of work generalizes or relaxes Shannon's notion to obtain unconditional guarantees under more realistic assumptions. The *wiretap channel* demonstrates that secrecy is achievable from channel asymmetries without pre-shared one-time keys (Wyner; Carleial–Hellman), with many extensions (e.g., Gaussian models) [12, 3, 5]. Another line shows how parties with correlated randomness can agree on secret keys through public discussion using *information reconciliation* and *privacy amplification* (Maurer; Bennett–Brassard–Robert; Bennett et al.; Maurer–Wolf), yielding unconditional security from weak common randomness [7, 2, 1, 8].

Closer to our operational viewpoint are *approximate* and *entropic* secrecy notions that quantify leakage rather than require $I(M; C) = 0$. Entropic security (Russell–Wang; Dodis–Smith) shows that if the message has sufficiently high min-entropy *from the adversary's perspective*, then short keys can suffice to ensure that no predicate of the message becomes easier to predict given the ciphertext [10, 4]. Subsequent work refines these guarantees and explores repetition and distribution-aware encryption [6]. These frameworks move from an all-or-nothing definition to quantitative bounds, which aligns with our Operational Perfect Secrecy (OPS) approach based on adversarial success probability.

Our contribution differs in two respects. First, we make explicit an *operational* ITS metric that directly upper-bounds the adversary's one-shot success probability by 2^{-t}, calibrated to deployment-level parameters (e.g., combinatorial richness of the extraction family). Second, we introduce two ITS forms tailored to systems with ephemeral public randomness blocks: *Combinatorial ITS (C-ITS)*, where ambiguity arises from the number of extractions consistent with a ciphertext and randomness, and *Dimensional Ambiguity ITS (DA-ITS)*, where the adversary cannot fix the probability space because both the key length and the extraction dimension are hidden. These notions complement entropic security and the wiretap/secret-key-agreement paradigms by providing clean, system-driven definitions and proofs that apply even when message distributions are structured and known.

8 Conclusion

This paper set out to resolve the fundamental gap between Shannon's theoretical notion of perfect secrecy and the requirements of real communication systems.

Shannon's framework established perfect secrecy as the strongest possible security notion, but it also imposed conditions that are rarely, if ever, satisfied in real systems. It assumes that messages are uniformly distributed, that keys are as long as messages, and that keys are drawn from an idealised source of perfect randomness. While mathematically consistent, this model is discontinuous: security collapses from full to none if any of these conditions are even slightly relaxed.

This paper has shown how to generalise Shannon's model into a form that is both conceptually and operationally usable. We introduced *Operational Perfect Secrecy (OPS)*, which defines secrecy as an explicit bound on adversarial success probability rather than as an all-or-nothing property of message posteriors. We proved that Shannon secrecy is recovered as the special case where the OPS parameter t equals the message length, while smaller t values give meaningful and quantifiable security guarantees.

Building on OPS, we defined two new families of information-theoretic security: *Combinatorial ITS (C-ITS)* and *Dimensional Ambiguity ITS (DA-ITS)*. C-ITS achieves OPS by creating large candidate sets of possible messages from short secret indices into ephemeral randomness, while DA-ITS adds uncertainty over the very dimensionality of the key space. Both constructions satisfy OPS and illustrate how information-theoretic security can be achieved with bounded key sizes.

In summary, this work reframes secrecy as a continuous, operational property, resolving the long-standing barrier created by Shannon's binary model. It provides a theoretical foundation on which scalable, deployable information-theoretic systems can be built, while strictly encompassing Shannon's definition as a limiting case.

Future Work. Future work will focus on implementing OPS-based security in practical systems and further analysing its composability properties. One direction is to develop concrete protocols that distribute short secret keys alongside large ephemeral randomness blocks, evaluating their performance and security empirically. Another is to formalise the interaction between OPS and existing security notions such as indistinguishability and semantic security, to clarify how OPS-based systems can integrate into existing cryptographic infrastructures. These steps will support the transition from theoretical foundations to deployable information-theoretic security architectures.

References

[1] Charles H. Bennett, Gilles Brassard, Claude Crépeau, and Ueli Maurer. Generalized privacy amplification. In *IEEE Transactions on Information Theory*, volume 41, pages 1915–1923, 1995.

[2] Charles H. Bennett, Gilles Brassard, and Jean-Marc Robert. Privacy amplification by public discussion. In *SIAM Journal on Computing*, volume 17, pages 210–229, 1988.

[3] A. B. Carleial and Martin E. Hellman. A note on wyner's wiretap channel. *IEEE Transactions on Information Theory*, 23(3):387–390, 1977.

[4] Yevgeniy Dodis and Adam Smith. Correcting errors without leaking partial information. In *TCC 2005*, volume 3378 of *LNCS*, pages 612–629. Springer, 2005.

[5] S. K. Leung-Yan-Cheong and Martin E. Hellman. The gaussian wiretap channel. *IEEE Transactions on Information Theory*, 24(4):451–456, 1978.

[6] Xiaoyu Li. Entropic security under repetition and structured message distributions. In *Information-Theoretic Cryptography (ITC)*, 2021.

[7] Ueli Maurer. Secret key agreement by public discussion from common information. *IEEE Transactions on Information Theory*, 39(3):733–742, 1993.

[8] Ueli Maurer and Stefan Wolf. Unconditionally secure key agreement and the intrinsic conditional information. *IEEE Transactions on Information Theory*, 45(2):499–514, 1999.

[9] Adrian Neal. Beyond shannon: Operational perfect secrecy as a generalised model for information-theoretic security, 2025.

[10] Alexander Russell and Hong Wang. How to fool an unbounded adversary with a short key. In *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 133–148. Springer, 2002.

[11] Claude E. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28(4):656–715, 1949.

[12] Aaron D. Wyner. The wire-tap channel. *Bell System Technical Journal*, 54(8):1355–1387, 1975.