

The General Epistemic Dispersion Theory

The Information Cost of Global Integrity in Fragmented Data

Massimo Comitato

Independent Researcher, Italy

Email: massimocomitato@gmail.com

ORCID: 0009-0002-8301-3589

The General Epistemic Dispersion Theory

Abstract

This work studies the informational cost of determining global integrity in fragmented data processed by distributed architectures. We distinguish between systems that require nodes to receive semantic and relational metadata to reconstruct the global state and systems in which global relations are encoded in a reusable structural core while peripheral nodes operate locally.

Under the stated informational assumptions, architectures lacking sufficient structure must transmit, pre-install, or reconstruct information needed to resolve residual configurational uncertainty. Conversely, when invariants, role associations, internal relations, and topology make the global configuration structurally determinable, the additional transcript may be empty or limited to residual structural uncertainty. Information is therefore shifted, not eliminated, from repeated communication to reusable structure.

A cost advantage arises when the amortized cost of building, distributing, accessing, and maintaining that structure is lower than the dynamic costs of metadata, coordination, semantic replication, and communication that it avoids.

Keywords: structural epistemic dispersion; information theory; distributed systems; global integrity; side information; structural information; conditional entropy.

Introduction

A general problem in the design of distributed systems concerns the information cost required to determine global integrity from a set of partial data.

A broad class of distributed architectures addresses this limitation by sharing among nodes the information required for semantic and relational understanding of the entire system state from the partial data. This approach entails a significant information cost, as it may require the repeated transmission, replication, storage, or reconstruction of explicit metadata to resolve local uncertainty.

This manuscript introduces General Epistemic Dispersion Theory, an analytical framework aimed at quantifying the distribution of the information cost required for the reconstruction and global validation of fragmented data. The framework formalizes the distinction between architectures based on explicit replication of semantic understanding and architectures in

which some of the relations required for global understanding are embedded in a reusable structural core.

By applying established principles of information theory—specifically, conditional entropy, source coding bounds, and the Fano inequality—the paper demonstrates how the topological structure of a system can function as implicit metadata. In the presence of structural sufficiency, some or all of the information required for configurational reconstruction or global integrity determination can be transferred from reiterated dynamic communication to a persistent and reusable structure.

This reallocation reduces the residual uncertainty that must be resolved with an additional transcript.

It can yield benefits in terms of communication, coordination, latency, and energy when the cost of building, distributing, accessing, and maintaining the structure is lower than the cost of the dynamic operations it avoids.

Original contribution and relation to prior work

The original contribution of this work lies in the formalization of structural epistemic dispersion as a general information-theoretic framework for the global validation of fragmented data. The framework explicitly decomposes the informational burden of global reconstruction between reusable structural information and residual dynamic information. It derives lower bounds on the amount of additional transcript required when structural sufficiency is incomplete, including the limiting case in which complete structural sufficiency renders the required additional transcript empty, i.e., $\ell(M_\Pi)=0$.

The individual mathematical tools employed in this work—including conditional entropy, the Fano inequality, source-coding bounds, and information complexity—are established results. It is also well known that the availability of side information, structural dependencies, or shared knowledge can reduce conditional entropy and, consequently, the amount of additional information required to reconstruct a state or compute a function.

The proposed innovation therefore does not lie in the introduction of new elementary principles of information theory, but rather in the formulation of a general framework in which the epistemic cost required for global understanding is explicitly reallocated from dynamic communication to a persistent and reusable structural core.

Positioning Relative to Major Prior Paradigms

Side information and correlated-source coding.

In the classical side-information paradigm, beginning with the results of Slepian and Wolf (1973), reductions in informational cost arise from the availability of information statistically correlated with the source to be encoded or reconstructed. In this setting, side information reduces residual uncertainty and, consequently, the amount of information required for reliable representation or reconstruction. In the present framework, the structural variable S may formally be interpreted as a form of side information, since its availability reduces $H(\Pi | F, S)$; however, this observation does not constitute the claimed novelty. The role assigned to S is more specific: it represents a persistent and reusable core composed of invariants, role associations, structural relations, feasibility constraints, and decoding geometry. The contribution therefore consists in formalizing such a structure as a carrier of configurational information capable of absorbing part of the uncertainty required for determining the global integrity of the raw data.

Functional coding (Orlitsky & Roche, 2001), graph-based side information, and index coding (Bar-Yossef et al., 2011).

In functional coding, graph-based side-information, and index-coding models, graphs and relational structures are employed to represent informational dependencies, knowledge already available to receivers, or constraints useful for reducing communication, coding, or computational cost. In the present framework, topology instead also assumes an epistemic and verification role: it does not merely describe information availability or exploitable dependencies, but contributes to determining which fragment configurations are compatible with a globally admissible state. Structural relations therefore become part of the integrity-determination mechanism, rather than serving exclusively as a means of optimizing information transmission or recovery.

Distributed snapshots and global-state determination.

Distributed snapshot algorithms, such as the Chandy–Lamport paradigm (Chandy & Lamport, 1985), address the problem of determining a consistent global state through a distributed procedure that records local states and communication channels. In such approaches, a global state is reconstructed through the explicit acquisition of information from the different components of the system. The present framework adopts a different separation of epistemic responsibilities: peripheral nodes may be limited to producing local observations or local results without having to possess, reconstruct, or interpret the entire global configuration. Integrity determination is delegated to the combination of local outputs and the persistent structural core S , thereby separating the local production of information from global state understanding.

Byzantine agreement and consensus.

Byzantine-agreement paradigms, and distributed consensus more generally (Lamport et al.,

1982; Dwork et al., 1988; Castro & Liskov, 1999), are oriented toward achieving a common decision among processes even in the presence of faults or adversarial behavior. In such architectures, communication, replication, and coordination are used to enable participants to converge toward a shared outcome. The objective of the present framework is different: global integrity determination does not require the nodes to reach agreement, nor does it require each participant to acquire a sufficiently complete representation of the overall state. Local knowledge may remain deliberately limited, while the ability to determine global validity is concentrated in the combination of local observations, structural relations, and internal invariants. The separation between local knowledge and global integrity determination therefore constitutes an architectural feature distinct from models based on distributed agreement.

Summary of the positioning.

The distinguishing feature of the framework therefore does not lie in the generic use of side information, graphs, or distributed information, but in the formalization of a persistent structural core as a reusable carrier of configurational information. This core reallocates part of the epistemic burden from dynamic communication to the structure of the system itself—its topology, relations, roles, constraints, and invariants—thereby reducing the amount of information that must be explicitly communicated or reconstructed during each execution. From this perspective, limited local knowledge does not necessarily constitute a deficiency of the system, but may instead become an architectural property compatible with the determination of global integrity.

The proposed innovation therefore lies in the formalization of structural epistemic dispersion as a reallocation of informational load between persistent structure and dynamic information in the problem of global validation of fragmented data.

Mathematical model

Let:

- X be the global raw state of an instance;
- $F = D(X)$ be the set of fragments generated by X ;
- $\mathcal{P}$ be the finite set of global feasible configurations;
- Π be a random variable that takes values in $\mathcal{P}$ and represents the realized global configuration of the fragmented instance, including roles, ordering, relationships, and instance-specific associations;
- $\mathcal{S}$ be a structural variable that takes values in the set $\mathcal{S}$ of feasible structural cores. Each realization s represents a reusable structural core containing feasibility constraints, invariants, structural relationships, role schemas, and decoding geometry. The same realization s can be reused across multiple compatible instances or

executions. By definition, S does not encode Π as a complete and execution-specific explicit description;

- M_Π be an additional configurational transcript possibly required for the complete reconstruction of Π ;
- $\mathcal{V}$ the deterministic global verification function;
- $Y \in \{\text{ACCEPT}, \text{VETO}\}$ the outcome of the global verification, defined by $Y = \mathcal{V}(F, \Pi, S)$
- M_Y an additional transcript possibly required to precisely determine outcome Y ;
- Π^* and Y^* , respectively, the reconstructed configuration and the reconstructed outcome of the global verification.

It is also assumed that fragmentation is informationally reversible once the correct configuration is known:

$$H(X \mid F, \Pi) = 0.$$

Equivalently, there exists a reconstruction function $\mathcal{R}$ such that:

$$X = \mathcal{R}(F, \Pi)$$

almost surely.

Configurational reconstruction and global verification

The complete configurational reconstruction is defined by:

$$\Pi^* = g_\Pi(F, S, M_\Pi) \quad (1),$$

with the exactness requirement:

$$P(\Pi^* = \Pi) = 1 \quad (2).$$

Equivalently:

$$P(\Pi^* \neq \Pi) = 0 \quad (3).$$

Since fragmentation is reversible with respect to the correct configuration, the exact reconstruction of Π also allows the lossless reconstruction of the global raw state:

$$X^* = \mathcal{R}(F, \Pi^*)$$

almost surely.

The outcome of the global verification can instead be reconstructed directly by:

$$Y^* = g_Y(F, S, M_Y) \quad (4),$$

with:

$$P(Y^* = Y) = 1.$$

Since Y is a deterministic function of F , Π , and S , the following holds:

$$H(Y \mid F, S) \leq H(\Pi \mid F, S) \quad (5).$$

Consequently, exact configurational reconstruction is sufficient to determine the outcome of the global verification, but it is not necessarily required in every architecture. A structural core may be sufficient to produce the exact outcome of the global verification even when it does not uniquely identify the entire implemented configuration.

Structural Sufficiency and Implicit Metadata

If a structural core S makes the global configuration Π determinable from fragments F , then the order, roles, relationships, and constraints do not need to be rewritten by an additional transcript at each run.

Thus, assuming that S is already available to the global reconstructor or validator, the additional runtime transcript needed to respecify Π can be empty:

$$M_{\Pi} = \emptyset \implies \ell(M_{\Pi}) = 0$$

This result does not imply that the construction, distribution, or maintenance of S is costless.

Theorem 1 — Residual informational cost of configurational reconstruction

The following theorem extends the previous lossless model to a general regime with bounded error; the exact case is recovered by setting $\varepsilon_{\Pi} = 0$.

Let Π be a discrete random variable that takes values in the finite set of admissible configurations $\mathcal{P}$, with $|\mathcal{P}| \geq 2$.

Let F be the set of observed fragments,

S be a structural variable that takes values in the set $\mathcal{S}$ of admissible structural cores.

Let

M_{Π} be an additional configurational transcript represented by a conditionally uniquely decodable binary encoding.

Define the residual structural uncertainty:

$$H_{\text{res}} := H(\Pi \mid F, S) \quad (6).$$

Let

$$\Pi^* = g_{\Pi}(F, S, M_{\Pi}) \quad (7)$$

be a reconstruction of the global configuration such that:

$$P(\Pi^* \neq \Pi) \leq \varepsilon_{\Pi} \quad (8)$$

with $0 \leq \varepsilon_{\Pi} \leq 1$,

and let

$p_{\Pi} := P(\Pi^* \neq \Pi)$ and therefore

$$p_{\Pi} \leq \varepsilon_{\Pi}.$$

Let us define δ_p the Fano informational term, which upper bounds the residual uncertainty compatible with an error probability ε_Π :

$$\delta_p(\varepsilon_\Pi) := h_2(\bar{\varepsilon}_\Pi) + \bar{\varepsilon}_\Pi \log_2 (|\mathcal{P}| - 1) \quad (9),$$

with

$$\bar{\varepsilon}_\Pi := \min \{ \varepsilon_\Pi, 1 - (1/|\mathcal{P}|) \} \quad (10).$$

Let

$$h_2(p) = -p \log_2 p - (1 - p) \log_2 (1 - p) \quad (11)$$

the binary entropy relating to the event

$$E = 1_{\{\Pi^* \neq \Pi\}},$$

i.e., error/no error.

Since the function

$$f(p) = h_2(p) + p \log_2 (|\mathcal{P}| - 1) \quad (12)$$

increases up to

$$p^* = 1 - (1/|\mathcal{P}|),$$

after this point, it decreases.

Therefore, the maximum of $f(p)$ for all actual error probabilities consistent with $p \leq \varepsilon_\Pi$ is obtained in $\bar{\varepsilon}_\Pi$.

Then the average length of the additional transcript satisfies:

$$\mathbb{E}[\ell(M_\Pi)] \geq \max\{0, H_{\text{res}} - \delta_{\mathcal{P}}(\varepsilon_\Pi)\} \quad (13).$$

Define

$$D_{\text{epi}}(S ; \Pi | F) := I(\Pi ; S | F) \quad (14)$$

and therefore:

$$D_{\text{epi}}(S ; \Pi | F) = H(\Pi | F) - H(\Pi | F, S) \quad (15).$$

The central decomposition follows:

$$H(\Pi | F) = D_{\text{epi}}(S ; \Pi | F) + H_{\text{res}} \quad (16).$$

Interpretation:

- $H(\Pi | F)$, initial configurational entropy;
- D_{epi} , part absorbed or reallocated in the structure;
- H_{res} , part that remains to be resolved dynamically.

Exact reconstruction

In case:

$$\varepsilon_\Pi = 0,$$

you have:

$$\delta_{\mathcal{P}}(0) = 0,$$

and therefore:

$$\mathbb{E}[\ell(M_\Pi)] \geq H_{\text{res}}.$$

Non-identifying regime in exact reconstruction

When:

$$H_{\text{res}} > 0,$$

the transcript cannot be identically empty and its average length must be strictly positive:

$$\mathbb{E}[\ell(M_\Pi)] \geq H_{\text{res}} > 0 \quad (17).$$

Identifying regime

When:

$$H_{\text{res}} = 0,$$

the configuration Π is almost surely determined by F and S .

There therefore exists a function g_0 such that:

$$\Pi = g_0(F, S) \text{ almost surely.}$$

In this case, the additional configurational transcript can be empty:

$M_\Pi = \emptyset$, and the reconstruction:

$$\Pi^* = g_0(F, S)$$

satisfies:

$$P(\Pi^* = \Pi) = 1, \ell(M_\Pi) = 0 \quad (18).$$

Proof

Since M_Π is represented by a conditionally uniquely decodable binary encoding and the reconstructor has F and S , the source coding theorem implies:

$$\mathbb{E}[\ell(M_\Pi)] \geq H(M_\Pi | F, S).$$

By the non-negativity of the conditional entropy:

$$H(M_\Pi | F, S) \geq I(\Pi ; M_\Pi | F, S),$$

since

$$I(\Pi ; M_\Pi | F, S) = H(M_\Pi | F, S) - H(M_\Pi | \Pi, F, S).$$

By definition of conditional mutual information:

$$I(\Pi ; M_\Pi | F, S) = H(\Pi | F, S) - H(\Pi | F, S, M_\Pi) \quad (19).$$

Since $\Pi^* = g_\Pi(F, S, M_\Pi)$, the reconstruction Π^* is a function of the information available to the reconstructor.

Therefore:

$$H(\Pi | F, S, M_\Pi) \leq H(\Pi | \Pi^*).$$

Applying the Fano inequality:

$$H(\Pi | \Pi^*) \leq h_2(p_\Pi) + p_\Pi \log_2 (|\mathcal{P}| - 1) \leq \delta_P(\varepsilon_\Pi) \quad (20).$$

Consequently:

$$H(\Pi | F, S, M_\Pi) \leq \delta_P(\varepsilon_\Pi).$$

Substituting:

$$I(\Pi ; M_\Pi | F, S) \geq H_{\text{res}} - \delta_P(\varepsilon_\Pi),$$

and therefore:

$$\mathbb{E}[\ell(M_\Pi)] \geq H_{\text{res}} - \delta_P(\varepsilon_\Pi) \quad (21).$$

Since an average length cannot be negative, the limit can be expressed as:

$$\mathbb{E}[\ell(M_\Pi)] \geq \max\{0, H_{\text{res}} - \delta_{\mathcal{P}}(\varepsilon_\Pi)\}.$$

In the exact case, $\varepsilon_\Pi = 0$, the Fano term is zero and we obtain:

$$\mathbb{E}[\ell(M_\Pi)] \geq H_{\text{res}} \quad (22).$$

Finally, when $H_{\text{res}} = 0$, the discrete nature of Π implies that the configuration is almost surely determined by F and S . Therefore, an exact reconstruction exists without additional configurational transcript.

Interpretation

Geometry is implicit, reusable metadata. The information doesn't disappear: it's contained in S , rather than being explicitly recommunicated during each reconstruction.

The epistemic weight of global understanding

Each node that must autonomously determine global integrity must receive, possess, or reconstruct enough information to resolve its uncertainty about the outcome of the global verification to handle partial data.

Let:

- $Y \in \{\text{ACCEPT}, \text{VETO}\}$, the outcome of the global verification;
- V_i , the local knowledge available to node i .

Let $Y_i^* = g_i(V_i, M_i)$ be an estimator of the global verification outcome,

with:

$$P(Y_i^* \neq Y) \leq \varepsilon.$$

If the node must calculate Y using a message M_i , with error at most ε , then:

$$I(Y; M_i | V_i) \geq H(Y | V_i) - h_2(\varepsilon),$$

with $0 \leq \varepsilon \leq 1/2$.

Consequently, assuming a uniquely decodable binary encoding of M_i :

$$\mathbb{E}[\ell(M_i)] \geq H(M_i | V_i) \geq I(Y; M_i | V_i) \quad (23).$$

Therefore:

$$\mathbb{E}[\ell(M_i)] \geq H(Y | V_i) - h_2(\varepsilon) \quad (24).$$

If, however, the node must only perform a local measurement:

$$L_i = f_i(V_i),$$

the global outcome is produced by the deterministic global verification function:

$Y = V_G(L_1, \dots, L_n; S)$.

Then the node does not need to receive the information needed to understand Y , Π , or the entire X .

Information complexity theory generally confirms that distributed computation of a function can require the revelation of a positive amount of information, even independently of the simple number of rounds (Kushilevitz & Nisan, 1997; Braverman & Rao, 2014).

Corollary — Residual Informational Cost of the Global Verification Outcome

If:

$$Y^* = g_Y(F, S, M_Y)$$

and

$$P(Y^* \neq Y) \leq \varepsilon_Y,$$

then, since Y is binary:

$$E[\ell(M_Y)] \geq \max\{0, H(Y | F, S) - h_2(\varepsilon_Y)\} \quad (25),$$

with

$$0 \leq \varepsilon_Y \leq 1/2.$$

In the exact regime:

$$\varepsilon_Y = 0,$$

we obtain:

$$E[\ell(M_Y)] \geq H(Y | F, S).$$

If

$$H(Y | F, S) = 0,$$

the outcome of the verification is already determined by F and S , even when:

$$H(\Pi | F, S) > 0.$$

Economic-structural Corollary

Note

All costs are expressed relative to a common cost function or mapped to a common metric, such as economic cost, normalized energy consumption, computational time, or an explicitly stated weighted combination.

Let C_S be the cost of building, deploying, and maintaining structure S .

Let C_E be the initial cost of the semantically explicit architecture.

Let $C_{local,t}^S$, and $C_{local,t}^E$ be the respective local costs in execution t , and let Δ_t be the additional cost of metadata, semantic replication, coordination, or communication required by the explicit architecture in execution t .

After Z executions:

$$C_{struct} = C_S + \sum_{t=1}^Z (C_{local,t}^S), \quad (26)$$

While, for the explicit architecture:

$$C_{\text{explicit}} = C_E + \sum_{t=1}^Z (C_{\text{local},t}^E + \Delta t) \quad (27).$$

The structure is overall less expensive when:

$$C_S - C_E + \sum_{t=1}^Z (C_{\text{local},t}^S - C_{\text{local},t}^E) < \sum_{t=1}^Z \Delta t. \quad (28)$$

In the simplified case where $C_E = 0$ and $C_{\text{local},t}^S = C_{\text{local},t}^E$,

the condition boils down to:

$$C_S < \sum_{t=1}^Z (\Delta t) \quad (29)$$

If the average savings per execution is $\Delta^* > 0$, the theoretical payback point is:

$$Z > C_S / \Delta^* \quad (30).$$

Interpretative limits

The results establish informational bounds conditional on the probabilistic distribution considered and the declared decodability assumptions. They do not automatically demonstrate a universal reduction in energy, latency, or cost in every single implementation. Such advantages require that the cost of building and maintaining the framework be lower than the cost associated with the explicit information and semantic replication thereby avoided.

Conclusion

Any structural uncertainty needed to determine integrity must be resolved through transmitted, pre-installed, or reconstructed information. A sufficient geometric structure transfers this information from repeated transcriptions to a reusable core, avoiding the need to replicate the global understanding at each local verifier.

Data Availability Statement

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

AI Assistance and Authorship Declaration

The Author acknowledges the use of AI tools for language translation and refinement, editorial revision, consistency checking, and assistance in mathematical formalization and verification of derivations. The conceptual architecture, research questions, theoretical framework, intended interpretation of the formal model, and scientific claims were conceived, directed, evaluated, and approved by the Author. The Author has reviewed the entire manuscript and assumes full scientific responsibility for its content.

Bibliography

- Aumann, R. J. (1976). Agreeing to disagree. *The Annals of Statistics*, 4(6), 1236-1239.
- Babai, L., Frankl, P., & Simon, J. (1986). Complexity classes in communication complexity theory. *Proceedings of the 27th Annual Symposium on Foundations of Computer Science (FOCS)*, 337-347.
- Bar-Yossef, Z., Birk, Y., Jayram, T. S., & Kol, T. (2011). Index coding with side information. *IEEE Transactions on Information Theory*, 57(3), 1479-1494.
- Bennett, C. H. (1982). The thermodynamics of computation—a review. *International Journal of Theoretical Physics*, 21(12), 905-940.
- Braverman, M., & Rao, A. (2014). Information Equals Amortized Communication. *IEEE Transactions on Information Theory*, 60(10), 6058-6069.
- Brewer, E. A. (2000). Towards robust distributed systems. *Proceedings of the 19th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, 7.
- Castro, M., & Liskov, B. (1999). Practical Byzantine Fault Tolerance. *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI)*, 173-186.
- Chandy, K. M., & Lamport, L. (1985). Distributed snapshots: Determining global states of a distributed system. *ACM Transactions on Computer Systems*, 3(1), 63-75.
- Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory* (2nd ed.). Wiley-Interscience.
- Dwork, C., Lynch, N., & Stockmeyer, L. (1988). Consensus in the presence of partial synchrony. *Journal of the ACM (JACM)*, 35(2), 288-323.
- Fagin, R., Halpern, J. Y., Moses, Y., & Vardi, M. Y. (1995). *Reasoning About Knowledge*. MIT Press.
- Fano, R. M. (1961). *Transmission of Information: A Statistical Theory of Communications*. MIT Press.
- Fischer, M. J., Lynch, N. A., & Paterson, M. S. (1985). Impossibility of distributed consensus with one faulty process. *Journal of the ACM (JACM)*, 32(2), 374-382.
- Gallager, R. G. (1968). *Information Theory and Reliable Communication*. John Wiley & Sons.
- Gilbert, S., & Lynch, N. (2002). Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. *ACM SIGACT News*, 33(2), 51-59.
- Halpern, J. Y., & Moses, Y. (1990). Knowledge and common knowledge in a distributed environment. *Journal of the ACM (JACM)*, 37(3), 549-587.
- Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review*, 106(4), 620-630.
- Kushilevitz, E., & Nisan, N. (1997). *Communication Complexity*. Cambridge University Press.
- Lamport, L., Shostak, R., & Pease, M. (1982). The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 4(3), 382-401.
- Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3), 183-191.
- MacKay, D. J. C. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press.

- Nash, J. (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1), 48-49.
- Nisan, N., Roughgarden, T., Tardos, E., & Vazirani, V. V. (2007). *Algorithmic Game Theory*. Cambridge University Press.
- Orlitsky, A., & Roche, J. R. (2001). Coding for computing. *IEEE Transactions on Information Theory*, 47(3), 903-917.
- Parrondo, J. M., Horowitz, J. M., & Sagawa, T. (2015). Thermodynamics of information. *Nature Physics*, 11(2), 131-139.
- Rényi, A. (1961). On measures of entropy and information. *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, 1, 547-561.
- Shamir, A. (1979). How to share a secret. *Communications of the ACM*, 22(11), 612-613.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(3), 379-423.
- Slepian, D., & Wolf, J. K. (1973). Noiseless coding of correlated information sources. *IEEE Transactions on Information Theory*, 19(4), 471-480.
- Szilard, L. (1929). On the decrease of entropy in a thermodynamic system by the intervention of intelligent beings. *Zeitschrift für Physik*, 53(11-12), 840-856.
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press.
- Yao, A. C. (1979). Some complexity questions related to distributive computing. *Proceedings of the 11th Annual ACM Symposium on Theory of Computing (STOC)*, 209-213.