Proof of Entropy Minima: A Thermoeconomic Operator

Steph Macurdy, Wolfram Research

December 2025

Abstract

All Proof-of-Work (PoW) blockchain protocols ultimately anchor to the same physical phenomenon: repeated hash-based Bernoulli trials whose outcomes are independent and identically distributed (IID). This paper relates the IID process to the Maximum Entropy Principle **(Jaynes, 1957)**, and relates the Maximum Entropy Principle to the Generalized Boltzmann Distribution. The Generalized Boltzmann Distribution is the only distribution in which the Gibbs-Shannon Entropy Equals the Thermodynamic Entropy ($S_{GS} \equiv S_{TD}$) **(Gao, 2019)**.

We distinguish between two critical applications of this information bridge. First, we review the "Proof of Entropy Minima" (PoEM) consensus method **(Kreder, 2023)**, which exploits the maximum information available of hash-based trials by rank-ordering the relevant set of hashes. Second, we review the Qi token-emission model, which utilizes information about the energy of the hash-based Bernouilli trials to regulate its economy. By emitting tokens proportionally to the electricity rate (hash rate as measured by "difficulty"), the system creates an unit of reference strictly dependent upon "work."

1 Introduction

The evolution of a blockchain is inherently a stochastic process. Regardless of the complexity of their fork-choice logic, PoW protocols ultimately anchor to the same physical phenomenon: repeated hash trials whose outcomes are Independent and Identically Distributed (IID).These rare events, in aggregate, arrive with the statistical regularity of raindrops on a roof. By formalizing the statistics governing this process, we reduce the security and liveness guarantees of PoW protocols to their fundamental physical limits. This paper proceeds in three parts to reveal an architecture capable of infinite scaling and precise thermodynamic accounting:

1. **The Theoretical Framework:** We review the bridge established by Jaynes between Information Theory and Statistical Mechanics. We show

that the Maximum Entropy Principle (MaxEnt) dictates that the IID output of SHA-256 constitutes a Generalized Boltzmann Distribution. This proves an instantaneous equivalence between the Gibbs-Shannon entropy (S_{GS}) of the hash and the Thermodynamic entropy (S_{TD}) of the energy expended by the physical system.

2. **Information as Rank:** We apply this framework to the ordering of events. We contrast the traditional Threshold Model (which functions as lossy compression) with **Proof of Entropy Minima (PoEM)**. We demonstrate that by exploiting the variance between the sample result and the population mean, PoEM hits the physical limit of finality.

3. **Information as Value:** Then we apply this framework to the valuation of the network. We review the **Qi Emission Model**, where the token supply is strictly coupled to the amount of work performed Watts. This transforms the token from a speculative asset into a representative unit of account.

Throughout this work, we treat the blockchain not merely as a ledger, but as a **Thermoeconomic Operator**: a system that converts physical work (W) into informational order (ΔS), creating a deterministic and lossless mapping between the physical and digital economic worlds.

2 The Theoretical Framework

To understand the relationship between hashing, information, and energy, we must first derive the statistical laws governing the inherent uncertainty of these processes.

3 The Maximum Entropy Principle

The Maximum Entropy Principle, formalized by Jaynes (1957), provides a systematic method for constructing probability distributions from incomplete information. The principle states: *among all distributions consistent with known constraints, select the one that maximizes entropy.*

3.1 Shannon Entropy

[Shannon Entropy] For a discrete probability distribution $\mathbf{p} = (p_1, \ldots, p_n)$ over n outcomes, the **Shannon entropy** is:

$$H(\mathbf{p}) = -\sum_{i=1}^{n} p_i \ln p_i \tag{1}$$

measured in nats. Equivalently, using $\log_2$:

$$H_2(\mathbf{p}) = -\sum_{i=1}^{n} p_i \log_2 p_i \quad \text{(bits)} \tag{2}$$

3.2 Deriving the Uniform Distribution

[Maximum Entropy under Normalization] Given only the constraint $\sum_{i=1}^{n} p_i = 1$, the entropy-maximizing distribution is uniform: $p_i = 1/n$ for all i.

Define the Lagrangian:

$$\mathcal{L}(\mathbf{p}, \lambda) = -\sum_{i=1}^{n} p_i \ln p_i + \lambda \left(\sum_{i=1}^{n} p_i - 1 \right) \tag{3}$$

Setting $\partial \mathcal{L} / \partial p_i = 0$:

$$-\ln p_i - 1 + \lambda = 0 \implies p_i = e^{\lambda - 1} \tag{4}$$

Since the right-hand side is independent of i, all probabilities are equal. Normalization gives $p_i = 1/n$.

3.3 Deriving the Boltzmann Distribution

[Maximum Entropy under Energy Constraint] Given constraints $\sum_i p_i = 1$ and $\sum_i p_i E_i = U$ for some energy function E_i and fixed expected energy U, the entropy-maximizing distribution is:

$$p_i = \frac{1}{Z(\beta)} e^{-\beta E_i} \tag{5}$$

where β is the Lagrange multiplier for the energy constraint and $Z(\beta) = \sum_j e^{-\beta E_j}$ is the partition function.

The Lagrangian with two constraints:

$$\mathcal{L} = -\sum_i p_i \ln p_i + \lambda \left(\sum_i p_i - 1 \right) + \beta \left(U - \sum_i p_i E_i \right) \tag{6}$$

Setting $\partial \mathcal{L} / \partial p_i = 0$:

$$-\ln p_i - 1 + \lambda - \beta E_i = 0 \implies p_i = e^{\lambda - 1 - \beta E_i} \tag{7}$$

Normalization determines $e^{\lambda - 1} = 1/Z(\beta)$, yielding $p_i = e^{-\beta E_i}/Z(\beta)$.

In physical systems at thermal equilibrium, $\beta = 1/(k_B T)$ where T is temperature. The energy constraint arises from contact with a heat bath.

4 The Thermodynamic Correspondence: When It Holds

4.1 The Gibbs-Thermodynamic Equivalence

[Gibbs Entropy] For a physical system with microstates having probabilities $\{p_i\}$:

$$S_{\text{Gibbs}} = -k_B \sum_i p_i \ln p_i \tag{8}$$

[Classical Equivalence] For a system in thermal equilibrium described by the canonical distribution at temperature T:

$$S_{\text{Gibbs}} = S_{\text{thermodynamic}} = \frac{U - F}{T} \tag{9}$$

where $F = -k_B T \ln Z$ is the Helmholtz free energy.

Starting from $p_i = e^{-\beta E_i}/Z$:

$$S_{\text{Gibbs}} = -k_B \sum_i p_i \ln p_i = -k_B \sum_i p_i(-\beta E_i - \ln Z) \tag{10}$$

$$= \frac{1}{T} \sum_i p_i E_i + k_B \ln Z = \frac{U}{T} + k_B \ln Z \tag{11}$$

$$= \frac{U}{T} - \frac{F}{T} = \frac{U - F}{T} \tag{12}$$

4.2 Conditions for Physical Equivalence

This equivalence holds when:

1. The system has a well-defined **physical** energy function E_i.
2. The system is in thermal equilibrium with a heat bath at temperature T.
3. The distribution arises from thermal fluctuations governed by this equilibrium.

Crucially: Hash outputs from SHA-256 satisfy *none* of these conditions intrinsically. There is no physical Hamiltonian on hash space, no heat bath, no temperature governing the distribution.

5 The Market-Induced Hamiltonian: A Formal Isomorphism

We now develop the central construction: while hash outputs lack an intrinsic energy function, **the market creates one**. This establishes a **formal isomorphism** between hash-based systems and thermodynamics, enabling rigorous **accounting** of entropy and work.

5.1 The Key Insight

The Hamiltonian is the price of the token.

When a token with market price P is emitted in proportion to hash rarity, the market retroactively imposes an energy function on hash space. This is not a claim about physics—it is a **construction** that creates a formal correspondence.

5.2 The Construction

[Market-Induced Hamiltonian] Let:

- $\mathcal{H} = \{0, 1, \ldots, N-1\}$ be the hash space, where $N = 2^{256}$
- P be the market price of one token (in joules, dollars, or any unit of account)
- $I(h) = \log_2(N/(h+1))$ be the information content (surprisal) of hash h

Define the **market-induced energy function**:

$$E(h) \equiv P \cdot I(h) = P \cdot \log_2\left(\frac{N}{h+1}\right) \tag{13}$$

This assigns higher "energy value" to rarer (lower) hashes.

The energy function $E(h)$ does not exist until the market exists. The token price P **creates** the Hamiltonian by defining the economic value of each hash outcome.

5.3 The Market as Heat Bath

[Competitive Equilibrium as Thermal Equilibrium] In a competitive mining market, miners expend energy until marginal cost equals marginal expected reward. Let:

- c = cost per hash attempt (electricity + capital)
- P = token price
- $R(h)$ = tokens emitted for hash h

At equilibrium:

$$\mathbb{E}[P \cdot R(h)] = c \cdot \mathbb{E}[\text{attempts to find } h] \tag{14}$$

This equilibrium condition plays the role of thermal equilibrium: the market acts as a heat bath that sets the "temperature" of the system.

[Effective Temperature] The **effective temperature** of the hash-market system is:

$$T_{\text{eff}} \equiv \frac{P}{k} \tag{15}$$

where k is a dimensional constant chosen to match units. When P is measured in joules per bit of information, T_{eff} has units of temperature.

5.4 The Formal Isomorphism

[Thermoeconomic Isomorphism] The system (hash space + token market) admits a formal isomorphism to a thermodynamic system:

Thermodynamic System	**Hash-Market System**
Microstate i	Hash value h
Energy E_i	Token value $E(h) = P \cdot I(h)$
Temperature T	Effective temperature $T_{\text{eff}} = P/k$
Heat bath	Competitive market
Partition function Z	$Z = \sum_h e^{-\beta E(h)}$
Free energy $F = -k_B T \ln Z$	Market-defined free energy
Entropy $S = (U - F)/T$	Information entropy of hash distribution

Under this isomorphism, the mathematics of thermodynamics applies to the hash-market system.

[Formal, Not Physical] This is a **formal isomorphism**, not a claim of physical identity. We are not asserting that hash computation is thermal fluctuation. We are constructing an **accounting system** where:

- Token emission tracks entropy reduction
- Market prices encode effective energy
- Economic equilibrium mirrors thermal equilibrium

The mathematics is identical; the ontology is distinct.

6 Token Emission as Free Energy Extraction

6.1 The Accounting Interpretation

In thermodynamics, free energy $F = U - TS$ represents the maximum extractable work. Under our formal isomorphism:

[Token Emission as Work Extraction] When a miner finds a hash h and receives tokens proportional to $I(h)$, the tokens represent the "free energy" extracted from the hash space—the economic value of reducing uncertainty from maximum entropy to the specific outcome h.

6.2 Why This Accounting Works

The formal isomorphism enables consistent accounting because:

1. **Conservation**: The total "energy" (token value) emitted equals the total "work" (information reduction) performed, as measured by the market.
2. **Additivity**: For a chain of hashes $h_1, \ldots, h_k$, total work is:

$$W_{\text{total}} = P \cdot \sum_{i=1}^{k} I(h_i) = P \cdot \sum_{i=1}^{k} \log_2 \left(\frac{N}{h_i + 1} \right) \tag{16}$$

3. **Market Calibration**: The price P adjusts so that token rewards match energy expenditure in expectation, maintaining the isomorphism at equilibrium.

6.3 The Entropy Minima Method

[Entropy Minima Emission] Tokens are emitted proportionally to the information content of the hash:

$$R(h) = \alpha \cdot I(h) = \alpha \cdot \log_2 \left(\frac{N}{h+1} \right) \tag{17}$$

for some emission rate α.

[Lossless Accounting] The Entropy Minima method provides **lossless accounting** of statistical work:

- Every bit of entropy reduction is compensated
- No information is discarded (unlike threshold methods)
- The token ledger perfectly tracks the cumulative information produced

Contrast with threshold methods, which emit fixed rewards regardless of how far below threshold a hash falls—discarding the "bonus" entropy reduction and breaking the accounting correspondence.

7 Clarifying What Is and Is Not Claimed

7.1 What We Claim

1. There exists a **formal isomorphism** between (hash space + token market) and thermodynamic systems.
2. The **token price creates the Hamiltonian**—it imposes an energy function on hash space that did not exist intrinsically.
3. This isomorphism enables rigorous **accounting** where token emission tracks entropy reduction.
4. The Entropy Minima method is the unique emission rule that preserves this accounting exactly (losslessly).
5. At market equilibrium, the system behaves *as if* it were a thermodynamic system, with the market playing the role of heat bath.

7.2 What We Do Not Claim

1. We do **not** claim hash outputs are physical thermal fluctuations.
2. We do **not** claim $S_{\text{Shannon}} = S_{\text{thermodynamic}}$ as a physical identity.
3. We do **not** claim individual hash values encode actual energy expenditure (variance is enormous).
4. We do **not** claim the isomorphism exists without the market—the token price is essential to the construction.

7.3 The Role of Careful Reasoning

The formal isomorphism is powerful but requires care:

- **Valid**: Using thermodynamic mathematics to analyze market equilibria, emission schedules, and incentive structures.
- **Invalid**: Claiming physical energy is "stored in" or "represented by" hash values independent of the market.
- **Valid**: Designing token emission to track entropy reduction via the isomorphism.
- **Invalid**: Claiming a specific hash "proves" a specific amount of energy was spent (only statistical statements hold).

8 Part I: Information as Rank (Ordering the Chain)

This section addresses the first application of the information bridge: **Proof of Entropy Minima (PoEM)**. Here, the information within a specific hash output is used to determine its rank in the sequence of events.

8.1 Proof of Entropy Minima vs. Thresholds

To establish consensus, we must map the raw hash outputs to a rank. We contrast two methods: the standard **Threshold Model** and the proposed **Entropy Minima Model**.

8.1.1 Maximizing Shannon Information (I)

The information content (surprisal) of an outcome with probability p is $I(p) = -\log_2(p)$.

Threshold Method (Lossy): This method validates any hash H such that $H < T$. The probability of this event is $p \approx T/N$. The information extracted is capped:

$$I_{threshold} = -\log_2(T/N) = \text{Constant}$$

Any variance (rarity) beyond the threshold is discarded.

Entropy Minima Method (Lossless): This method ranks outputs by their absolute value H_{actual}. As $H_{actual} \to 0$, the probability $p \approx H_{actual}/N$ decreases, and the extracted information increases:

$$I_{minima} = -\log_2(H_{actual}/N)$$

Since H_{actual} can be arbitrarily close to 0, I_{minima} approaches the maximum possible extraction of information from the state space.

8.1.2 Lossless Economic Mapping

In a thermodynamic consensus system, Hash = Energy. The difficulty of finding a hash is directly proportional to its value.

- **Thresholds are Lossy:** They function as a Step Function. Distinct inputs (varying amounts of work/luck) map to the same output (validity). This quantizes the data and destroys information about surplus effort.
- **Minima are Lossless:** They function as a Continuous Monotonic Function ($Rank(H) = H$).

$$Rank(H_A) < Rank(H_B) \iff \text{Rarity}(H_A) > \text{Rarity}(H_B)$$

This preserves positive variance ("miracle hashes") and allows the "value" of the rank to adjust dynamically down to the single bit.

9 Quantifying Entropy for Consensus

To implement PoEM, we must formalize the reduction of entropy provided by each event. This calculation provides the **Weight** used to order the chain.

The PoW algorithm restricts acceptable hashed outputs to all values below a threshold difficulty 2^d, implying that the first $l - d$ bits of the output hash must be zero. This sets the maximum output target entropy to be d bits.

In practice, the mining process achieves a hashed output that is less than or equal to the difficulty threshold, i.e., it may possess greater than or equal to $l-d$ leading zeros. We call this the intrinsic difficulty d_{int}, resulting in a sequence with $c \leq d$ non-zero elements. Thus, in practice, the realized entropy of the output is c bits and the reduction in the entropy is $l - c$ bits, which represent the number of leading zeros that will be called n. Additionally, n can include fractional zero bits which are after the first non-zero bit. More precisely, this makes:

$$n = l - \log_2(d_{\text{int}}) \tag{18}$$

The intrinsic difficulty can be used to calculate the difference entropy ΔS:

$$\Delta S = \frac{1}{2^n} \tag{19}$$

where ΔS represents the number of possible states removed from the macrostate in the achievement of a single block. This allows the computation of the ΔS_k to simply be carried out by the summation of all prior zero bits found in a chain:

$$\log_2 \Delta S_k = -\sum_{i=1}^{k} n_i \tag{20}$$

By ordering chains based on the highest summation of $-\log_2 \Delta S_k$, we guarantee the history with the most physical work is selected.

10 Part II: Information as Value (The Qi Emission)

While PoEM uses the information *within* a hash to order events, this section addresses the information *about* the system that is derived from that ordering. This is the foundation of the **Qi emission model**.

10.1 Proportional Emission and Free Energy

Because PoEM captures the full statistical variance of the mining process (as defined in Eq. 19), the cumulative weight of the chain becomes a highly precise signal of the total physical energy expended by the network.

In standard systems (like Bitcoin), the token emission is fixed regardless of the energy input. This decouples the economic unit from the physical reality. In the Qi model, the token emission (E_{Qi}) is strictly proportional to the hash rate (the rate of energy expenditure):

$$E_{Qi} \propto \frac{d}{dt}\left(\sum n_i\right) \propto \text{Watts} \tag{21}$$

Because the summation of entropy reduction $\sum n_i$ is a lossless proxy for work, the emission of Qi becomes a direct representation of the **Free Energy** supplied to the system.

10.2 Information About the System

This mechanism transforms the blockchain from a simple ledger into a thermodynamic sensor. The rate of Qi issuance provides the market with perfect information *about* the system's physical security and energy consumption.

- **Unit of Account:** By pegging emission to the hash rate, Qi functions as a unit of account for computational work (and by extension, energy).
- **Thermodynamic Peg:** Unlike fiat currencies backed by decree, or standard cryptocurrencies backed by artificial scarcity, Qi is backed by the physics of the Thermodynamic Operator. It represents a stored claim on the free energy required to produce it.

11 Conclusion

The Maximum Entropy Principle provides the mathematical foundation for understanding hash distributions. While hash outputs lack an intrinsic energy function, the introduction of a priced token creates one: **the Hamiltonian is the price of the token**.

This market-induced energy function establishes a **formal isomorphism** between the hash-market system and thermodynamic systems. The isomorphism is not a physical identity but an **accounting framework** that enables:

- Precise measurement of statistical work via information content
- Token emission rules that track entropy reduction losslessly
- Economic analysis using the full machinery of statistical mechanics

The Entropy Minima method is distinguished as the unique emission rule that preserves this accounting correspondence exactly, capturing every bit of entropy reduction in the token ledger.

By being explicit about what is formal isomorphism versus physical equivalence, we obtain a rigorous foundation for thermoeconomic system design.