

Collapse as Thermodynamic Optimization: A Logical Framework for Deriving Born's Rule

Manas Mandhan

April 30, 2025

Abstract

This paper presents a thermodynamic route to the Born rule, proposing that quantum wavefunction collapse is a physical process governed by informational entropy minimization. Using Landauer's principle, we treat collapse as the erasure of alternatives in a superposed quantum state, and assign a thermodynamic cost to this erasure. We show that the collapse outcome distribution which minimizes the Kullback-Leibler divergence from the state's internal informational structure corresponds uniquely to the Born rule. This approach reframes quantum probabilities as entropy-optimal and collapse as a real, irreversible process constrained by information-theoretic and thermodynamic principles.

1 Introduction

In the standard formalism of quantum mechanics, measurement is treated as an exception. The evolution of a quantum system is governed by unitary dynamics until an observation is made. Then, according to Born's rule, the system's wavefunction collapses, and one result is realized with a probability proportional to the squared amplitude of that state.

But this rule is not derived from deeper principles. It is inserted into the theory as a postulate. The collapse process is not explained, only prescribed. This leads to a foundational question: is there a physical reason the Born rule holds?

The measurement problem has inspired a wide range of interpretations in quantum mechanics, each attempting to resolve how or why a definite outcome emerges from a superposition. In the Copenhagen interpretation, collapse is triggered by observation, but remains fundamentally undefined. Many-worlds avoids collapse entirely by positing that all outcomes occur in branching universes. Objective collapse theories introduce stochastic dynamics to

force reduction. QBism treats probability as a subjective belief. Decoherence explains the appearance of classicality, but does not explain how a single result is selected.

These approaches diverge in metaphysics but share a common structure: none of them derive the Born rule from first principles. Instead, it is either inserted axiomatically or assumed as a statistical regularity. In this paper, I take a different route: I assume that collapse occurs and that it carries an entropy cost. From this assumption, I explore whether the Born rule can be derived as the optimal solution to an entropy minimization problem.

My central idea is that collapse is a thermodynamic process, one constrained by the physical cost of erasing uncertainty. Applying Landauer’s principle—the idea that erasing information incurs a minimum entropy cost—I propose that collapse favors outcomes that minimize information loss [1, 2].

This interpretation does not rely on subjective observation, external branching, or hidden variables. It frames collapse as a real, physical event governed by thermodynamic efficiency, rather than metaphysical necessity or mathematical formalism. The aim is not to discard existing interpretations but to reframe the question itself: If collapse has a cost, what is the most efficient way for it to occur?

The goal of this paper is to show that the Born rule may arise naturally as the least-cost solution to this problem—a distribution that minimizes entropy divergence from the pre-collapse state. I make no assumptions about decoherence, hidden variables, or even subjective probability. Only for the case that collapse occurs, costs something, and is resolved efficiently. The rest, as I argue, follows logically.

2 Motivational Foundations

2.1 Why Treat Collapse as Thermodynamic

Collapse, if real, must be physically meaningful. It reduces a quantum superposition into a single outcome and discards the rest. That process erases information. And as Landauer’s principle tells us, information erasure has a thermodynamic cost [1].

Landauer’s principle, initially framed for classical information, has been extended into the quantum regime and shown to hold under general conditions [2]. This gives us a clear physical lens: collapse must incur an entropy cost for the information it eliminates.

From this, I take a second step: if collapse carries a cost, nature may prefer the most efficient version of it.

2.2 A Principle of Efficiency

Efficiency here means minimizing the amount of informational entropy lost in the collapse. Rather than collapsing randomly, the system may collapse in a way that optimizes the cost of erasure. This reframes collapse not as a mystery, but as a thermodynamic optimization problem.

Specifically, I argue that the final outcome distribution—the probabilities we associate with the outcomes—is the result of minimizing the informational distance between the pre-collapse and post-collapse states.

To formalize this optimization, I use Shannon entropy [3] to represent uncertainty in the system and Kullback-Leibler (KL) divergence [4] to define the cost of transitioning from one probability distribution to another.

KL divergence quantifies how much information is "lost" when approximating one distribution with another. If collapse transitions from the superposed amplitudes to a definitive outcome, KL divergence gives us a measure of how costly that change is in informational terms.

2.3 Why I Do Not Use Decoherence (Yet)

In a follow-up paper, I intend to incorporate decoherence and Hamiltonian dynamics into this model to explore the environmental and temporal evolution of collapse. But here, I begin from first principles, with no assumptions about external observers or branching realities. Just collapse, cost, and the logic of optimal erasure.

3 Entropy Framework for Collapse

3.1 Shannon Entropy in a Quantum System

Consider a pure quantum state written in an orthonormal basis $\{|i\rangle\}$ corresponding to the eigenstates of a measurable observable:

$$|\psi\rangle = \sum_i c_i |i\rangle, \quad \text{where } c_i \in \mathbb{C}, \quad \sum_i |c_i|^2 = 1.$$

At this stage, I do not interpret $|c_i|^2$ as physical probabilities. Rather, I treat them as structural weights associated with the decomposition of the state. I do not assume the Born rule, which would equate these squared amplitudes with measurement probabilities.

Suppose that collapse must occur, and that the system ultimately transitions to a definite outcome $|i\rangle$. I define a *candidate* probability distribution over these outcomes:

$$P = \{p_1, p_2, \dots, p_n\}, \quad \text{where } p_i \in [0, 1], \quad \sum_i p_i = 1.$$

This distribution represents the uncertainty associated with the collapse process. Importantly, p_i is not yet derived—it is a free parameter over which the collapse mechanism will eventually optimize.

I define the Shannon entropy of this candidate distribution as:

$$H(P) = - \sum_i p_i \log p_i.$$

This entropy quantifies the uncertainty inherent in the collapse outcome *before* a result is realized. Collapse to a definite state implies that this entropy must be fully erased.

Thus, in this framework, $H(P)$ captures the informational structure that must be destroyed during collapse. The central idea is that this erasure carries a thermodynamic cost, and that collapse proceeds in a way that minimizes the divergence from the initial informational structure of the state.

3.2 Landauer’s Principle and the Cost of Erasure

Landauer’s principle establishes a deep connection between information theory and thermodynamics. It states that the erasure of one bit of information from a physical system necessarily incurs a minimal entropy cost. In its classical form, the principle is written as:

$$\Delta Q \geq k_B T \ln 2,$$

where ΔQ is the heat dissipated into the environment, k_B is the Boltzmann constant, and T is the temperature of the thermal reservoir.

Originally formulated for classical computing systems by Rolf Landauer in 1961, the principle was motivated by the realization that only logically irreversible operations—such as erasing a bit—are accompanied by physical irreversibility. More recently, Landauer’s bound has been extended to quantum systems, where it applies to the erasure of quantum information and the resetting of quantum memories.

In the context of quantum measurement, collapse can be viewed as an irreversible operation that discards all but one outcome from a superposition. This process reduces uncertainty, and thus, according to Landauer’s principle, must generate a corresponding entropy

cost.

If a quantum system transitions from a state containing multiple possible outcomes to a single definite outcome, it is performing a physical act of information erasure. It eliminates the alternative outcomes, rendering them physically unrealized. Landauer’s principle then implies that this erasure process cannot be thermodynamically free.

In this paper, I treat wavefunction collapse as such an act of erasure. My central assumption is that if collapse incurs a thermodynamic cost, then the system has reason to collapse in a way that minimizes that cost.

This motivates the search for a cost function—an expression that captures the entropy increase associated with collapsing one informational structure into another. In the next section, I argue that this cost is naturally quantified by the Kullback-Leibler (KL) divergence between the final outcome distribution and the internal informational structure of the system.

3.3 Minimizing Divergence to Select the Collapse Distribution

If collapse incurs a thermodynamic cost, then nature has reason to prefer outcome distributions that minimize it. The goal is now to determine which probability distribution $\{p_i\}$ over outcomes minimizes the information cost of collapsing from a superposed quantum state.

I define this cost using the Kullback-Leibler (KL) divergence:

$$D_{\text{KL}}(P \parallel Q) = \sum_i p_i \log \left(\frac{p_i}{q_i} \right),$$

where $Q = \{q_i\}$ is the system’s internal reference distribution, and $P = \{p_i\}$ is the collapse outcome distribution.

At this point, we ask: what should Q be?

The Elimination of Arbitrary Distributions

Suppose we try a uniform reference distribution, $q_i = 1/N$, or assign equal weights to all outcomes regardless of the system’s internal state. This would imply that the system forgets its structure entirely before collapse—a claim clearly violated by decades of quantum experiments. Real-world collapse outcomes show a consistent statistical alignment with the squared amplitudes of the quantum state. Uniform or arbitrary distributions do not match what is observed.

Collapse is not random. It reflects something internal to the quantum system—and that something must relate to how the system interacts with its environment.

Environment as Informational Constraint

The only viable candidate for Q must emerge from the interaction between the system and its environment. This is precisely the insight decoherence theory provides. While I do not formally invoke decoherence in this paper, I acknowledge its explanatory power.

In decoherence, when a system entangles with an environment and is partially traced over, the reduced density matrix becomes diagonal in the pointer basis. The diagonal elements—interpreted as effective probabilities—are precisely the squared amplitudes $|c_i|^2$ of the system state in that basis.

This suggests that the system’s interaction with the environment naturally encodes a distribution:

$$q_i = |c_i|^2.$$

Born Rule as the Entropy-Minimizing Solution

Now, substituting $q_i = |c_i|^2$ into the KL divergence and minimizing the cost:

$$\min_P \left[\sum_i p_i \log \left(\frac{p_i}{|c_i|^2} \right) \right],$$

subject to $\sum_i p_i = 1$, yields the minimum when:

$$p_i = |c_i|^2.$$

Thus, the Born rule emerges as the ****only collapse distribution**** that minimizes thermodynamic cost while remaining consistent with the system’s environment-constrained informational structure.

4 Interpretations, Implications, and Assumptions

The result of this framework is conceptually simple but interpretively deep: the Born rule emerges not as a postulate, but as the unique outcome of a thermodynamically efficient collapse process. If collapse is real, and if it carries an entropy cost in the sense outlined by Landauer, then minimizing that cost selects the outcome distribution $p_i = |c_i|^2$. This reframes the Born rule as a principle of informational economy.

Relation to Quantum Interpretations

This approach is compatible with—but distinct from—existing interpretations of quantum mechanics. It does not require the metaphysical branching of Many-Worlds, nor the subjective Bayesianism of QBism. It sidesteps the vagueness of Copenhagen’s measurement postulate and offers a concrete physical mechanism that objective collapse theories often leave undefined.

Importantly, this interpretation does not reject decoherence but repositions it. Decoherence explains how classicality emerges from entanglement with the environment. My proposal suggests that the probabilities themselves—long treated as axiomatic—may have their origin in thermodynamic regularities. If decoherence supplies the structure of $q_i = |c_i|^2$, then the Born rule emerges as the entropy-optimal resolution of collapse.

Implications

This framework introduces a new lens for viewing quantum probabilities. Rather than treating them as fundamental or observer-dependent, I treat them as constrained by the information content of the system itself. Collapse is modeled not as a metaphysical discontinuity, but as a process governed by real-world principles: cost, erasure, and efficiency.

The implication is that quantum probabilities may be ****emergent****, not fundamental—derived from deeper physical laws governing how uncertainty must resolve under thermodynamic constraints.

Assumptions and Scope

To arrive at this conclusion, several key assumptions were made:

- **Collapse is real and irreversible.** It is not an illusion or a branching of worlds, but a process that erases alternatives.
- **Erasures has a thermodynamic cost.** Landauer’s principle is assumed to hold at the level of quantum state resolution.
- **The system carries an internal informational structure.** This is expressed as a reference distribution q_i , later identified with $|c_i|^2$.
- **No unitary dynamics or decoherence were included.** The framework is purely informational and thermodynamic in its reasoning.

These assumptions are deliberately minimal, and are not intended to fully explain the quantum-to-classical transition. Rather, the goal of this work is to show that even within this stripped-down framework, Born’s rule arises naturally and uniquely.

5 Future Work: From Logical Framework to Physical Derivation

The assumptions laid out in the previous section—regarding collapse, thermodynamic cost, and internal informational structure—were introduced as minimal yet necessary premises to derive the Born rule as a least-cost outcome. These were not postulated for metaphysical comfort, but chosen to match a physically plausible and conceptually clean mechanism for measurement.

In upcoming work, I intend to move beyond logical clarity and toward formal physical grounding. The goal is to show that these assumptions are not free-standing philosophical choices, but natural consequences of established quantum thermodynamics, decoherence theory, and system-environment dynamics.

Specifically, I aim to:

- **Collapse and irreversibility:** Model quantum measurement as a logically irreversible process, represented by tracing over system-environment interactions in an open Hamiltonian framework. This connects the collapse process to the general structure of entropy production as formalized in quantum thermodynamics [2].
- **Thermodynamic cost of erasure:** Use the improved quantum Landauer bound by Reeb and Wolf [5] to quantify the entropy cost of erasing alternatives during collapse. This will allow the cost function used here—based on KL divergence—to emerge from first principles.
- **Derivation of the reference distribution q_i :** Employ decoherence theory to show how system-environment entanglement leads to a reduced density matrix whose diagonal elements correspond to $|c_i|^2$ in the pointer basis. This provides the formal structure needed to justify the use of $q_i = |c_i|^2$ as the reference distribution in the entropy minimization argument.
- **Embedding the optimization in quantum dynamics:** Extend the KL-based variational framework into a time-dependent setting, using the tools of open system Hamil-

tonians and resource-theoretic thermodynamics to model the evolution of collapse as a physical transition constrained by energy, entropy, and information flow.

This future work will transform the current proposal from a thermodynamically inspired interpretation into a physically grounded mechanism. If successful, it will show that the Born rule is not a postulate, but a necessary consequence of quantum irreversibility, information erasure, and environmental structure.

References

- [1] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM journal of research and development*, 5(3):183–191, 1961.
- [2] Massimiliano Esposito and Christian Van den Broeck. Second law and landauer principle far from equilibrium. *EPL (Europhysics Letters)*, 95(4):40004, 2011.
- [3] Claude E. Shannon. A mathematical theory of communication. *Bell system technical journal*, 27(3):379–423, 1948.
- [4] Solomon Kullback and Richard A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22(1):79–86, 1951.
- [5] David Reeb and Michael M. Wolf. An improved landauer principle with finite-size corrections. *New Journal of Physics*, 16(10):103011, 2014.