# Relational Probability Fields: Coherence, Collapse and Cross-Domain Patterns

**Veronika Pudsey**

Independent Researcher

human-quietspace@proton.me

## Abstract

Why do such different systems—quantum measurements, thermal ensembles, Bayesian updates and modern language models—keep producing probability mappings with strikingly similar structure? We introduce relational probability fields: distributions over discrete possibilities shaped by context-dependent relational potentials (Hamiltonians, logits, likelihoods, fitness) and resolved by collapse or flow. We show that the Born rule, Boltzmann statistics and softmax sampling (and their Bayesian and evolutionary counterparts) can be expressed within a single template: potentials induce an outcome distribution, and resolution instantiates particular outcomes. We further introduce interface style descriptors—concentration, context sensitivity and memory—to characterise how sharply and how stably this resolution behaves across domains, and to operationalise a system's characteristic relational mode of turning context into realised outcomes.

Building on this shared structure, we propose a speculative interpretation: probabilities can be read as measures of relational coherence (fit) between configurations and their context, and resolution can be viewed as a form of coherence optimisation on the space of possibilities. In this view, Born weights $|\langle\lambda\_i|\psi\rangle|^2$ quantify geometric fit in Hilbert space, while Boltzmann factors and softmax scores play analogous roles in energy and representation spaces. We argue that making this coherence structure explicit may clarify why familiar probabilistic laws take the forms they do, and may connect naturally to information-theoretic and variational approaches to gravity, where dynamics is driven by extremal principles on relational structures.

The paper is primarily conceptual: it does not propose a new microphysical theory, nor does it derive the Born rule or Einstein's equations from a single principle. Instead, it offers a compact cross-domain language for comparing probabilistic mechanisms and a coherence-based template for generating testable questions at the interface of quantum theory, statistical mechanics, machine learning and gravitational physics.

# 1 Introduction

The Born rule is one of the most distinctive features of quantum mechanics. It postulates that if a system is described by a normalized state $|\psi\rangle$ and an observable $A$ with discrete spectrum is measured, then the probability of obtaining eigenvalue $\lambda_i$ is given by

$$P(\lambda_i) = \langle\psi \mid P_i \mid \psi\rangle,$$

where $P_i$ projects onto the eigenspace of $\lambda_i$. For one–dimensional eigenspaces this reduces to

$$P(\lambda_i) = |\langle\lambda_i \mid \psi\rangle|^2.$$

Born first introduced this statistical interpretation in 1926 in the context of scattering theory, and it was later formalized and discussed in detail by von Neumann and many others. (Born, 1926; Landsman, 2009; Neumaier, 2019)

Despite its central role, the Born rule is usually not derived from more elementary assumptions. It is accepted as a rule that connects a complex-valued field (the wavefunction) to measurement statistics. This raises a natural question:

Why this mapping from an underlying field to probabilities, and not another?

A similar question arises in other domains. In equilibrium statistical mechanics, probabilities over microstates $\sigma$ are given by a Boltzmann distribution,

$$P(\sigma) = \frac{e^{-\beta H(\sigma)}}{Z},$$

where $H(\sigma)$ is the energy of configuration $\sigma$, $\beta$ is inverse temperature and $Z$ a normalizing partition function. (Baxter, 1982) In modern machine learning, large language models compute a vector of logits $z$, then convert it into token probabilities using a softmax function

$$P_i = \frac{e^{z_i}}{\sum_j e^{z_j}},$$

an approach that can be traced back to log-linear and neural network models for classification and sequence modeling. (Bridle, 1990; Vaswani et al., 2017; Radford et al., 2019)

These three rules—Born, Boltzmann and softmax—look different and operate on different quantities (complex amplitudes, energies, real-valued scores). But they share a recognizable pattern:

1. There is a **field of possibilities** (eigenvalues of an observable, spin configurations, possible next tokens).

2. There is a **relational potential** that assigns each possibility a weight depending on its relations to other variables and to a context (Hamiltonian, interaction couplings, logit scores).
3. There is a **mapping from potentials to probabilities** that determines how likely each possibility is to be realized when a selection or collapse occurs.

In quantum theory, collapse is associated with a measurement event. (Landsman, 2009; Neumaier, 2019) In statistical mechanics, individual microstates are sampled from the equilibrium distribution by physical dynamics or Monte-Carlo algorithms. (Baxter, 1982) In language models, each token generation step samples a symbol from the model's probability distribution as defined by a Transformer-style architecture. (Vaswani et al., 2017; Radford et al., 2019)

In each case, a structured probability field is resolved into realised outcomes, either via discrete selections or via continuous flows.

The purpose of this paper is to make this shared structure explicit and to formulate a minimal, substrate-agnostic framework for relational probability fields, their resolution mechanisms (collapse or flow), and the interface style by which resolution behaves across contexts. We show that:

- quantum measurement with the Born rule,
- equilibrium ensembles in Ising-like systems with Boltzmann statistics, and
- token sampling in large language models with softmax logits

can all be viewed as instances of a common pattern: a probability field shaped by relational potentials, and a dynamics that resolves this field into realized states.

We also name a system's relational mode—the characteristic way it turns context into realised outcomes—operationalised via interface style descriptors.

We do **not** claim that quantum systems, statistical ensembles and language models are physically equivalent, nor that all of them are "quantum" in any strong sense. The claim is more modest and more structural: once we shift from an object-based description ("a particle has a property") to a field-based relational description ("a configuration has a weight in a context"), these domains can be described using a shared vocabulary.

In the final part of the paper we point out that similar field-and-collapse mechanisms appear in Bayesian inference and in evolutionary game dynamics, where distributions over hypotheses or types are shaped by data or fitness landscapes. (Jaynes, 2003; Hofbauer & Sigmund, 1998, 2003) Taken together, these examples motivate the question of whether there exists a more general principle according to which relational potentials over fields of possibilities determine the probabilities of realized outcomes.

The paper proceeds as follows. Section 2 introduces the abstract language of relational probability fields, relational potentials, mappings to probabilities, and resolution mechanisms, and defines interface style descriptors (concentration, context sensitivity, and memory). Section 3 applies this language to three concrete systems: quantum measurement, Boltzmann ensembles and large language models. Section 4 briefly extends the framework to Bayesian inference and evolutionary game dynamics. Section 5 sketches a speculative coherence-based interpretation of the resulting probability fields, including a possible bridge toward gravitational dynamics. Section 6 discusses scope, limitations and connections to existing work, and Section 7 concludes.

# 2 Relational probability fields: potentials, resolution and interface style

In what follows we separate three roles that are often conflated in domain-specific formalisms. First, a potential $\Phi$(or an amplitude field, energy function, score, likelihood, etc.) defines a structured field over possibilities. Second, a mapping rule specifies how this structure is converted into an outcome distribution $P(\cdot \mid C)$. Third, a resolution mechanism realises particular outcomes (as discrete collapses or as continuous flows), and can be characterised by an interface style: how strongly probability mass concentrates, how sensitively it responds to contextual perturbations, and whether effective outcome statistics exhibit memory (history dependence). In addition, we define relational mode as the signature way these roles are coupled in a given domain, and we use interface style as a practical comparative handle on it. These distinctions will be used throughout the paper to compare systems without committing to a common microdynamics.

## 2.1 Probability fields

Let $\mathcal{S} = \{s_1, \dots, s_K\}$be a discrete set of possible outcomes, states or configurations. A probability field over $\mathcal{S}$is a probability distribution

$$P: \mathcal{S} \to [0,1], \sum_{i=1}^{K} P(s_i) = 1.$$

In many physical and computational systems, this distribution is not arbitrary. It is induced by a lower-level relational structure:

- in quantum mechanics, by a complex amplitudes field,
- in statistical mechanics, by an energy or Hamiltonian on configurations,
- in machine learning, by a vector of scores or logits produced by a model.

We denote this lower-level structure by a potential $\Phi(s_i; C)$, where $C$represents a context: external parameters, interactions, boundary conditions, previous tokens, and so on. The potential encodes how each possibility relates to other components of the system and to its environment.

## 2.2 Relational potentials

The central idea is that $\Phi$ is not an intrinsic “property” of a state $s_i$, but a function of relations:

- between degrees of freedom within the state (e.g. spin–spin couplings),

- between the state and external fields or apparatus,
- or between the state and a history or prompt.

Schematically, we write

$$\Phi(s_i; C) = \Phi(s_i \mid \text{relations to other variables in } C).$$

In an Ising model, for example, the energy $H(\sigma)$of a spin configuration $\sigma$is given by a sum over pairwise couplings and external fields,

$$H(\sigma) = -\sum_{\langle i,j \rangle} J_{ij}\, \sigma_i \sigma_j - \sum_i h_i\, \sigma_i,$$

where the sum $\langle i, j \rangle$runs over neighboring sites, $J_{ij}$are pairwise couplings and $h_i$local external fields. (Baxter, 1982) Here $H$is a relational potential over configurations.

In a language model, logits $z_i$depend on learned weights connecting the current hidden state to possible next tokens; these weights capture statistical relations in the training data and the local context represented by a prompt. (Vaswani et al., 2017; Radford et al., 2019)

In a quantum system, amplitudes $\psi_i$for different eigenstates of an observable arise from the system's Hamiltonian and its prior evolution; interference patterns reflect phase relations between components of the wavefunction. (Born, 1926; Landsman, 2009; Neumaier, 2019)

## 2.3 Mapping potentials to probabilities

Different systems use different rules to map potentials to probabilities. In this paper we focus on three canonical cases:

1. **Born rule (quantum measurement).**
   For an orthonormal eigenbasis $\{\mid \lambda_i \rangle \}$, amplitudes $\alpha_i = \langle \lambda_i \mid \psi \rangle$define probabilities

   $$P(\lambda_i) = \mid \alpha_i \mid^2.$$

2. **Boltzmann distribution (statistical mechanics).**
   For an energy function $H(\sigma)$, equilibrium probabilities are

   $$P(\sigma) = Ze - \beta H(\sigma), Z = \sigma' \sum e - \beta H(\sigma').$$

3. **Softmax distribution (large language models).**
   For logits $z_i \in \mathbb{R}$, token probabilities are

   $$P_i = \frac{e^{z_i}}{\sum_j e^{z_j}}.$$

Boltzmann weights and softmax probabilities both take an explicit normalized exponential form,

$$P(s_i \mid C) = \frac{e^{\phi_i(C)}}{\sum_j e^{\phi_j(C)}},$$

where $\phi_i(C)$ is a real-valued score induced by the underlying relational potential (for Boltzmann, $\phi_i = -\beta H(s_i)$; for softmax, $\phi_i = z_i$).

The Born rule is not an exponential-family mapping in this sense: the underlying quantities $\alpha_i$are complex amplitudes that support interference and whose probabilistic interpretation is constrained by Hilbert-space geometry. Nevertheless, it plays the same functional role in the relational template: it converts a structured field (here, amplitudes projected onto a measurement basis) into a probability field over realised outcomes.

This distinction matters for substance but not for the organisational pattern emphasised in this paper. In the next subsection we introduce interface style descriptors—properties such as concentration (sharpness), context sensitivity and memory—that characterise how probability fields behave under contextual change and resolution, without conflating the different microphysical mechanisms that generate them.

## 2.4 Interface style: shape, sensitivity and memory

So far we have treated the mapping from relational potentials to probabilities as a "rule" (Born, Boltzmann, softmax) and treated collapse/selection as a black-box step that draws a realised outcome from the resulting probability field. This is sufficient to make the shared triad explicit. However, if the goal is to use the relational template as a comparative tool across domains, we need a way to characterise how an interface maps fields of possibilities into probability fields, and how those probability fields are resolved into outcomes in practice.

We therefore introduce the notion of interface style. By "interface" we mean the effective relational boundary at which a probability field is instantiated and resolved: a measurement arrangement in quantum mechanics, a thermal environment and coupling in statistical mechanics, a decoding rule and sampling procedure in language models, or an update rule in Bayesian and evolutionary settings. By "style" we mean a small set of cross-domain descriptors that characterise the shape of the induced probability field and its response to contextual changes, without committing to any particular microdynamics.

Formally, given a set of possibilities $S = \{s_1, \ldots, s_K\}$ and context $C$, the relational mechanism specifies a map

$$(\Phi, C) \mapsto P(\cdot|C),$$

and a realisation rule that resolves $P(\cdot \mid C)$into a particular outcome $s \in S$(either by sampling, by deterministic selection, or by continuous flow). The style of an interface is the family of properties of these maps that remain meaningful across domains.

Below we highlight three minimal style descriptors that will recur throughout the examples.

### 2.4.1 Concentration (sharpness) of the probability field

The first descriptor is the concentration (or sharpness) of the probability field $P(\cdot | C)$: whether probability mass is distributed broadly across many possibilities or concentrated on a small subset.

A simple proxy is the Shannon entropy

$$H(P) = -\sum_{i=1}^{K} P(s_i \mid C) \log P(s_i \mid C),$$

or, equivalently, the effective support size $\exp(H(P))$. (Shannon, 1948) Low entropy corresponds to a sharp field (few dominant alternatives); high entropy corresponds to a diffuse field (many comparable alternatives).

This descriptor is explicitly tunable in several systems. In Boltzmann ensembles, inverse temperature $\beta$controls concentration: low temperature (high $\beta$) concentrates weight near low-energy configurations, while high temperature spreads weight more uniformly across configuration space. In language models, decoding temperature plays an analogous role: decreasing temperature sharpens the distribution over tokens, increasing temperature diffuses it. In quantum measurement, concentration depends on the prepared state and on the measurement context: some system–measurement configurations yield near-deterministic outcome statistics, while others yield broad distributions over eigenvalues.

Concentration does not by itself specify which outcomes are favoured. It characterises the shape of the induced field: how decisively an interface "selects" among its available possibilities.

### 2.4.2 Context sensitivity (threshold behaviour)

The second descriptor is context sensitivity: how strongly the probability field changes when the context is perturbed. Intuitively, some interfaces are robust, producing similar probability fields under small changes in $C$; others are near-threshold, where small contextual changes can cause large redistribution of probability mass.

One way to capture this is through a divergence between fields at neighbouring contexts, for example the total variation distance or the Kullback–Leibler divergence: (Kullback & Leibler, 1951)

$$D(P(\cdot \mid C), P(\cdot \mid C')),$$

for contexts $C'$ that differ from $C$ by a small perturbation. High sensitivity corresponds to large divergence under small contextual change; low sensitivity corresponds to small divergence.

This descriptor is particularly important for comparing systems at regime boundaries. In statistical mechanics, sensitivity increases near critical points where small changes in external fields or couplings can lead to large shifts in macroscopic behaviour. In language models, small prompt changes can move the system between different continuation regimes, reshaping the token distribution. In quantum settings, outcome statistics can depend sharply on measurement basis choice, coupling strength, or on the structure of system–environment interactions in open systems.

In Bayesian inference, sensitivity reflects how strongly new data reweights hypotheses relative to the prior.

In the relational framework, context sensitivity captures a key feature of “collapse-like” resolution: not merely that outcomes are drawn probabilistically, but that the distribution of possible outcomes can be highly contingent on contextual relations.

### 2.4.3 Interface memory (Markov vs non-Markov resolution)

The third descriptor concerns whether the mapping and resolution are effectively memoryless or whether they carry history dependence. In many idealised formulations, the probability field $P(\cdot|C)$depends only on the present context $C$. But in many real systems the effective interface has internal degrees of freedom that retain information about past interactions, so that outcome statistics depend on history in ways not captured by the instantaneous context alone.

We use interface memory as a neutral descriptor for this distinction. At one extreme, a Markov-like interface is fully specified by the present context $C$ (up to the chosen model). At the other extreme, a non-Markov interface requires an expanded context $C^{*}$ that includes relevant history variables, internal states, or environmental records.

Interface memory is well studied in open quantum systems, where non-Markovian environments can feed information back into a system and change its effective measurement statistics. (Breuer & Petruccione, 2002) It also appears in classical contexts as hysteresis, path dependence and slow variables in coupled dynamics. In language models, memory is explicit in the token context window, but additional forms of history dependence can also be induced by the decoding procedure and by external interaction loops. In evolutionary and market-like systems, memory appears as persistent structures that modulate future selection pressures.

In our framework, interface memory matters because it changes what it means for a probability field to be “context-dependent”: the context may not be a static list of external parameters but a dynamically evolving relational structure.

### 2.4.4 Style as a cross-domain comparative layer

These descriptors—concentration, context sensitivity and interface memory—do not replace the detailed mathematics of any domain. Rather, they form a comparative layer that can be applied to any instance of the relational triad (field, potential, collapse/flow). Two systems may share the same abstract template while exhibiting very different interface styles: one may generate sharp, low-sensitivity distributions with little memory; another may generate diffuse, high-sensitivity distributions with strong history dependence.

Introducing style descriptors serves two purposes. First, it makes the relational template more informative: we can compare not only that different domains instantiate the same triad, but also how their probability fields behave under contextual change. Second, it provides a natural bridge to information-theoretic and geometric tools—entropy, divergences, Fisher-type metrics—that quantify these properties and may support future empirical or computational studies.

These style descriptors are not new physical postulates; they repackage familiar notions—entropy-based sharpness, response/susceptibility, and Markov vs non-Markov dependence—into a cross-domain comparative layer. (Shannon, 1948; Kullback & Leibler, 1951; Breuer & Petruccione, 2002)

In the remainder of the paper we will not attempt a systematic measurement of style across domains. We will, however, occasionally point out where particular examples are naturally associated with sharpness control (e.g., temperature), threshold sensitivity (e.g., critical regimes) or interface memory (e.g., open systems and history-dependent updates), as these are precisely the features that become visible once probability is treated as a relational field rather than as an intrinsic property of isolated states.

### 2.4.5 Relational mode

The shared template (field–potential–mapping–resolution) describes what components a probabilistic mechanism has. Interface style descriptors describe how the induced probability field behaves (sharpness, sensitivity, memory). Together, these point to a deeper comparative notion: systems can share the same template while differing in the characteristic way they relate context, potentials and realised outcomes. We call this the relational mode of an interface: the system's distinctive way of turning contextual relations into a distribution and then into realised records. Relational mode is not an additional layer on top of interface style. It is the name for what interface style descriptors collectively measure: the characteristic way a system relates context to realised outcomes. Concentration, sensitivity and memory are not separate parameters — they are observable projections of relational mode, the handles by which it becomes empirically accessible across domains.

## 2.5 Collapse and selection

The final ingredient is a collapse or selection mechanism: a rule or process that turns a probability field into a realized state.

- In quantum mechanics, a measurement of observable $A$yields one eigenvalue $\lambda_i$with probability $P(\lambda_i)$, and the post-measurement state is typically modeled as the corresponding eigenstate $|\lambda_i\rangle$. (Landsman, 2009; Neumaier, 2019)
- In statistical mechanics, physical dynamics or Monte-Carlo sampling produce specific microstates according to the Boltzmann distribution. (Baxter, 1982)
- In large language models, token sampling procedures (greedy decoding, temperature sampling, nucleus sampling, etc.) select one token at each step from the model's probability distribution. (Vaswani et al., 2017; Radford et al., 2019)

In all three cases, the selection is irreversible at the level of realized outcomes: once a specific eigenvalue, configuration or token has been instantiated, the prior superposition or distribution is no longer accessible as such. One can reconstruct or re-estimate a probability field from many runs, but individual collapses are one-way events.

We will use "collapse" as a neutral term for this transition from a structured field of possibilities to a specific outcome. The key claim of this paper is that, across the three systems we analyze in detail and several others we mention briefly, collapse events have a common structure: they are realizations drawn from probability fields shaped by relational potentials.

# 3 Three systems

In this section we consider three concrete cases and rewrite them in the common language of “field – potential – collapse”:

1. quantum measurement (Born rule),
2. the Ising/Boltzmann model in statistical physics,
3. next-token prediction in a large language model.

## 3.1 Quantum measurement as relational collapse

Consider a quantum system prepared in a normalized state $|\psi\rangle$ in a finite-dimensional Hilbert space, and an observable $A$ with discrete, non-degenerate spectrum,

$$A = \sum_{i=1}^{K} \lambda_i \mid \lambda_i\rangle\langle\lambda_i \mid.$$

Before measurement, the state can be expanded in the eigenbasis of $A$:

$$\mid \psi\rangle = \sum_{i=1}^{K} \alpha_i \mid \lambda_i\rangle, \alpha_i = \langle\lambda_i \mid \psi\rangle, \sum_i \mid \alpha_i \mid^2 = 1.$$

Here the probability field over possible outcomes $\{\lambda_i\}$is given by the Born rule:

$$P(\lambda_i \mid \psi, A) = \mid \alpha_i \mid^2.$$

In terms of interface style (Section 2.4), the concentration of this field ranges from near-deterministic (when $|\psi\rangle$ is close to an eigenstate of $A$) to highly diffuse (when amplitudes are broadly distributed in the measurement basis). Sensitivity is immediate: changing the measurement context (e.g., the chosen observable or basis) generically reshapes the entire distribution. In more realistic settings, system–environment coupling can also introduce interface memory, so that effective outcome statistics depend not only on instantaneous context but also on history.

The underlying relational potential resides in the complex amplitudes $\alpha_i$, which are determined by the prior evolution of the system under its Hamiltonian and by its preparation. Phases and interference encode relations between components: in a different basis, the same state $|\psi\rangle$generally has different amplitude patterns. (Born, 1926; Landsman, 2009; Neumaier, 2019)

When a measurement of $A$ is performed, an outcome $\lambda_i$ is realized with probability $P(\lambda_i)$, and the post-measurement state is typically modeled as the corresponding eigenstate $|\lambda_i\rangle$ (or, more generally, as the projection of $|\psi\rangle$ onto the eigenspace associated with $\lambda_i$). (Landsman, 2009; Neumaier, 2019)

In the language of this paper:

- the **field of possibilities** is the set of eigenvalues $\{\lambda_i\}$,
- the **relational potential** is encoded in the amplitude pattern $\{\alpha_i\}$, built from relations between the system, its Hamiltonian and its preparation context,
- the **collapse** is the measurement event's selection of one eigenvalue according to the Born probabilities.

We do not attempt to reinterpret the measurement problem or to derive the Born rule. The point is more modest: viewed structurally, quantum measurement takes a relationally structured field (the wavefunction projected into an eigenbasis) and resolves it into a specific outcome via a probabilistic rule.

## 3.2 Ising / Boltzmann ensembles as relational fields

We now turn to a prototypical model in statistical mechanics: the Ising model. (Baxter, 1982) Consider a set $\Lambda$ of lattice sites, each carrying a spin variable $\sigma_i \in \{-1, +1\}$. A configuration $\sigma = (\sigma_i)_{i\in\Lambda}$ assigns a spin value to each site.

The energy (Hamiltonian) of a configuration is

$$H(\sigma) = -\sum_{\langle i,j\rangle} J_{ij}\ \sigma_i\sigma_j - \sum_i h_i\,\sigma_i,$$

where the sum $\langle i, j\rangle$ runs over neighboring sites, $J_{ij}$ are pairwise couplings and $h_i$ local external fields.

At inverse temperature $\beta = 1/(k_B T)$, the equilibrium probability of configuration $\sigma$ is given by the Boltzmann distribution

$$P(\sigma \mid \beta, J, h) = \frac{e^{-\beta H(\sigma)}}{Z(\beta, J, h)}, Z = \sum_{\sigma'} e^{-\beta H(\sigma')}.$$

The interface style is particularly transparent here. Concentration is controlled by temperature: increasing $\beta$ (lowering $T$) sharpens the field around low-energy configurations, while decreasing $\beta$ diffuses it. Context sensitivity can become large near critical regimes, where small changes in couplings or external fields produce disproportionate redistribution of probability mass. Finally, while the equilibrium distribution itself is memoryless, the realised sampling dynamics can exhibit history dependence through slow variables, metastability or hysteresis in non-equilibrium settings.

In our framework:

- the **probability field** is the distribution $P(\sigma)$ over all $2^{|\Lambda|}$ possible spin configurations;

- the **relational potential** is the energy function $H(\sigma)$, which is purely relational: it depends on pairwise products $\sigma_i\sigma_j$and on relations to external fields $h_i$, not on any spin in isolation;

- the **collapse / selection** corresponds to the system occupying a particular configuration $\sigma$at a given time or, algorithmically, to a single sample drawn from the Boltzmann distribution.

The Ising model is particularly transparent as a relational field: the coupling matrix $J_{ij}$encodes how strongly spins prefer to align or anti-align with their neighbors. Configurations where these preferences are jointly satisfied have lower energy and therefore higher probability in equilibrium. (Baxter, 1982)

From the point of view of probability fields, the Boltzmann distribution implements a softmax over negative energy:

$$P(\sigma) = \frac{e^{-\beta H(\sigma)}}{\sum_{\sigma'} e^{-\beta H(\sigma')}}.$$

The energy landscape defined by $H$determines which regions of configuration space are more likely to be realized when the system relaxes or when we sample from the distribution.

## 3.3 Large language models as relational probability fields

Consider a large language model (LLM) trained to predict the next token in a sequence, typically implemented as a Transformer network (Vaswani et al., 2017). Given a context $C$(a sequence of previous tokens), the model computes a hidden representation $h(C)$and then a vector of logits $z(C) \in \mathbb{R}^K$, where $K$is the vocabulary size. The probability of choosing token $t_i$as the next symbol is given by a softmax function (Bridle, 1990; Radford et al., 2019),

$$P(t_i \mid C) = \frac{e^{z_i(C)}}{\sum_{j=1}^{K} e^{z_j(C)}}.$$

In terms of interface style, concentration is explicitly tunable by decoding parameters such as temperature (and, indirectly, by truncation schemes such as top-$k$ or nucleus sampling), which reshape how probability mass is distributed over tokens. Context sensitivity is often high: small prompt changes can shift logits and redistribute probabilities across the vocabulary. Interface memory is implemented by the evolving context window (the token history that defines $C$); in interactive settings, additional memory-like effects can arise from external loops that feed previous outputs back as new context.

In these terms we can identify the same three components as before:

- the **field of possibilities** is the set of candidate next tokens $\{t_1, \dots, t_K\}$,

- the **relational potential** is given by the logits $z_i(C)$, which depend on learned weights and on the specific context $C$,

- the **collapse / selection** is the sampling step that chooses one token according to $P(\cdot \mid C)$.

The logits themselves are not arbitrary scores. They are the output of a deep neural network whose parameters encode statistical regularities of text and other training data (Radford et al., 2019). For a simple prompt such as

The meaning of life is

a trained model may assign relatively high logits to continuations such as “to”, “not”, “a”, “that”, “unknown”, “something”, “just”, “life”, “love”, and very low logits to most other tokens. After applying softmax, this yields a high-entropy probability field over possible next tokens, with a modest preference for a small set of semantically plausible continuations and tiny probabilities spread across the rest of the vocabulary.

When the model samples a specific token—say, “to”—the probabilistic field over all possible continuations is resolved into one realized continuation. The context is then updated by appending this token, a new hidden state $h(C')$is computed for the extended context $C'$, and a new probability field over next tokens is generated. This process repeats step by step.

Structurally, this is very close to the Boltzmann case in statistical mechanics: the logits $z_i(C)$play the role of a context-dependent potential, and the softmax function plays the role of a normalized exponential mapping from potentials to probabilities (Bridle, 1990). In both systems, a relationally defined scalar function over a discrete set of possibilities is turned into a probability field, from which individual realizations are drawn, with interface style determining how sharply and how sensitively this resolution behaves.

# 4 Further instances of the same pattern

The three systems considered so far—quantum measurement, Ising/Boltzmann ensembles and large language models—make the common structure particularly transparent: in each case we have an explicitly defined probability field over a discrete set of possibilities, a relational potential that shapes this field, and a context-dependent collapse or selection mechanism.

In this section we briefly point to two additional domains where an analogous structure appears:

1. Bayesian inference, where a prior field over hypotheses is reshaped by a likelihood.
2. Evolutionary game dynamics, where a distribution over types is driven by a fitness landscape.

We do not analyze these systems in the same level of detail, but we include them to indicate that the pattern is not limited to quantum theory, statistical mechanics and machine learning.

## 4.1 Bayesian inference as probabilistic field update

In Bayesian inference, one starts from a prior distribution $P(H)$over a discrete or continuous set of hypotheses $H$, and updates this distribution in light of data $D$. The posterior distribution is given by Bayes’ rule,

$$P(H \mid D) = \frac{P(D \mid H)\, P(H)}{\sum_{H'} P(D \mid H')\, P(H')},$$

where $P(D \mid H)$is the likelihood function (Jaynes, 2003).

In the language of this paper:

- the **probability field** is the prior distribution $P(H)$over hypotheses,
- the **relational potential** is encoded in the likelihood $P(D \mid H)$, which expresses how strongly each hypothesis relates to, or is supported by, the observed data $D$,
- the **collapse / selection** can be interpreted in two ways:
  - as the update from prior to posterior, which reshapes the entire field, or
  - as the choice of a single hypothesis, for example via maximum a posteriori (MAP) estimation or sampling from the posterior.

Formally, Bayes' rule has the multiplicative structure

$$P_{\text{posterior}}(H) \propto W(H; D)\, P_{\text{prior}}(H),$$

where $W(H; D) = P(D \mid H)$plays the role of a **data-dependent relational weight**. Up to normalization, the new probability field is obtained by reweighting the old field by a context-dependent factor that favours hypotheses better aligned with the data.

This is structurally similar to the exponential reweighting in Boltzmann distributions and softmax mappings. In all three cases, a pre-existing probability field is reshaped by a context-dependent relational potential (energy, score, likelihood), and realized outcomes are then drawn from the updated field.

## 4.2 Evolutionary game dynamics as flow on a fitness landscape

In evolutionary game dynamics, one studies how the frequencies of different types or strategies change over time under the influence of selection and interaction (Hofbauer & Sigmund, 1998, 2003). Let $x_i(t)$denote the relative frequency of type $i$at time $t$, with $\sum_i x_i(t) = 1$. The classic replicator equation is

$$\dot{x}_i = x_i(f_i(x) - \bar{f}(x)),$$

where $f_i(x)$is the fitness (or payoff) of type $i$in population state $x$, and $\bar{f}(x) = \sum_j x_j f_j(x)$is the average fitness.

Here we can identify:

- the **probability-like field** as the distribution $x(t) = (x_1(t), \ldots, x_K(t))$over types,
- the **relational potential** as the fitness values $f_i(x)$, which typically depend on interactions between types (e.g., via a payoff matrix or game structure),

- the **collapse / selection** not as a discrete sampling event, but as a flow of the field $x(t)$toward attractors in the fitness landscape.

Fitness is relational in a direct sense: in evolutionary games, the payoff to a strategy depends on how it interacts with the strategies played by others, so that $f_i(x)$is determined by the configuration of the whole population (Hofbauer & Sigmund, 1998). The replicator dynamics then reweights the distribution over types in a way that amplifies those with above-average fitness and suppresses those with below-average fitness, causing the population state $x(t)$to move "uphill" on the fitness landscape toward fixed points, cycles, or more complex attractors.

From the perspective of this paper, the key point is that the population state defines a probability-like field over types, the fitness landscape defines a relational potential over that field, and the evolutionary dynamics implements a context-dependent reweighting akin to the probabilistic collapses we have seen in other systems. The difference is that selection is here realized as a continuous flow rather than as a sequence of discrete sampling events, but the underlying pattern—field shaped by relations, resolved by a context-dependent rule—remains present.

## 4.3 Summary

Bayesian inference and evolutionary game dynamics are conceptually very different domains. Yet, once we describe them in terms of fields over possibilities and relational potentials that reshape those fields, they can be placed alongside quantum measurement, Boltzmann ensembles and large language models.

In all five cases we find:

- a distribution (or distribution-like state) over possibilities,
- a context-dependent relational structure (likelihood, fitness, Hamiltonian, network scores),
- and a rule that uses this structure to determine which possibilities are amplified, suppressed or realized.

These further instances suggest that the triad of probability field – relational potential – collapse or flow is not an artefact of a particular formalism, but a recurring pattern across domains that are usually treated as unrelated.

# 5 Collapse as coherence optimisation: a speculative hypothesis

Sections 2–4 established a structural template: fields of possibilities shaped by relational potentials, mapped to outcome distributions and resolved by collapse or flow. In this section we do not add new structure. Instead, we propose a speculative interpretation of what the weights in these distributions might represent dynamically. The structural claims of the paper do not depend on accepting this interpretation.

Up to this point we have treated collapse or selection as a black-box mechanism: a process that draws realised outcomes from a probability field shaped by a relational potential. In this section we briefly sketch a more speculative line of thought that attempts to interpret these probabilities in dynamical terms.

The key idea is to treat probabilities not merely as summaries of frequencies or states of knowledge, but as measures of relational coherence or attraction strength between configurations and their context.

## 5.1 A generic coherence functional

In several of the systems considered in this paper, probabilities arise by normalized reweighting of a discrete set of possibilities by a context-dependent score. A common instance is a normalized exponential form,

$$P(s_i \mid C) \propto \exp(\phi_i(C)),$$

where $\phi_i(C)$is a real-valued function induced by an underlying relational potential. Concretely:

- in Boltzmann ensembles, $\phi_i = -\beta H(s_i)$;
- in large language models, $\phi_i = z_i(C)$(the logit of token $t_i$in context C);
- in Bayesian updating, $\phi(H) = \log P(D \mid H) + \log P(H)$;
- in evolutionary and game-theoretic settings, $\phi_i$can be taken as a fitness or payoff-based score (e.g. $f_i(x)$) when selection is modelled by discrete choice rules (logit/softmax) or related approximations.

Quantum measurement differs in its mathematical structure: the Born rule is not a normalized exponential mapping over real-valued scores. Instead, probabilities are given by squared projection amplitudes $P_i = |\alpha_i|^2$, where $\alpha_i = \langle \lambda_i \mid \psi \rangle$are complex and support interference. Nevertheless, at the level relevant for the present conceptual move, Born weights still define an ordering and a relative weighting over outcomes that can be related to a score by taking $\phi_i = \log |\alpha_i|^2$as an analogue.

At a purely formal level we may therefore define a relational coherence functional

$$\mathcal{C}(s_i; C) := \phi_i(C),$$

such that higher $\mathcal{C}$implies higher probability. In this language, the probability field

$$P(s_i \mid C) \propto \exp(\mathcal{C}(s_i; C))$$

can be read as saying: configurations with greater relational coherence with respect to the current context $C$exert a stronger “attraction” on the system.

The interface style introduced in Section 2.4 then characterises how this coherence score is transferred into an outcome distribution and realised outcomes—e.g., how sharply probability mass concentrates, how sensitively it responds to contextual perturbations, and whether the effective interface carries memory.

By “coherence” we do not mean quantum phase coherence in the narrow technical sense, but a more general notion of relational fit or compatibility:

- in an Ising or Boltzmann ensemble, low-energy configurations are those in which local couplings and external fields are jointly satisfied; they exhibit a high degree of fit to the interaction structure encoded in $H$;
- in a language model, high-logit tokens are those that fit the semantic, syntactic and statistical regularities captured in the network weights and the prompt context $C$;
- in Bayesian inference, hypotheses with large $\log\ P(D \mid H) + \log\ P(H)$are those that jointly fit prior structure and observed data.

In quantum theory, the natural measure of “fit” between a state $\mid \psi\rangle$and an eigenstate $\mid \lambda_i\rangle$of an observable is the squared inner product $\mid \langle\lambda_i \mid \psi\rangle \mid^2$. This quantity is singled out by the geometry of Hilbert space and by the requirement that probabilities assigned to an orthonormal decomposition sum to one. On the present reading, Born weights $\mid \alpha_i \mid^2$can be seen as measuring how strongly the system–apparatus configuration coheres with each eigenstate of the measured observable. The Born rule then appears not merely as an ad hoc postulate, but as a natural way of quantifying relational coherence in a Hilbert space setting.

## 5.2 Collapse as resolution toward coherent configurations

If we adopt this perspective, collapse or selection can be viewed, at least tentatively, as a process of coherence optimisation:

- In statistical mechanics, physical dynamics causes the system to explore configuration space and to spend most of its time in regions where $-\beta H(\sigma)$is large, that is, where energy is low and $C$is high.
- In large language models, the decoding procedure preferentially selects tokens with higher logits in a given context, i.e., tokens that achieve higher semantic coherence with the prompt and with learned relational structure.
- In Bayesian inference, the posterior reweights hypotheses according to how well they cohere with the data and prior structure, via $\log\ P(D \mid H) + \log\ P(H)$.

In all these cases, configurations with higher $C$are effectively more “attractive” in the space of possibilities: they are more likely to be realised, more stable under perturbations, or more favoured by the dynamics. This does not eliminate stochasticity. Rather, it reframes it: randomness remains in the choice of which configuration is realised on a given run, but the probability law governing these choices is interpreted as encoding relational coherence. Instead of asking why individual outcomes are random, we ask why randomness follows this particular distribution. The answer, on this view, is that $P(s_i \mid C)$ quantifies how strongly each configuration fits into the relational structure defined by the system and its context.

Importantly, the coherence-based reading does not by itself determine whether resolution is near-deterministic or highly stochastic: this depends on interface style (concentration and sensitivity) as well as on the specific realisation regime (sampling, thermalisation, projective measurement, etc.).

A speculative extension is to apply the same intuition to quantum measurement. Instead of treating Born probabilities as primitive, one could ask whether $|\alpha_i|^2$might be interpreted as a measure of quantum relational coherence between the system, the eigenstate $|\lambda_i\rangle$and the measurement apparatus. Collapse would then not be a completely featureless random jump, but the resolution of a relationally defined field toward configurations of higher coherence, subject to constraints we do not yet understand.

At this stage, this is only a conceptual suggestion. A full account would need to specify:

- what, precisely, quantum coherence $C_{QM}$is in terms of amplitudes, phases and system–apparatus interactions;
- how such a coherence functional relates to the standard unitary dynamics and to the Born rule;
- and whether collapse should be viewed as a fundamentally stochastic process guided by $C$, or as an effective description of an underlying deterministic but inaccessible dynamics.

We do not attempt to answer these questions here; we only note that, structurally, the pieces are in place for such a reinterpretation.

## 5.3 Possible connections to gravity

If probabilities are read as measures of relational coherence, it becomes natural to ask whether some macroscopic "effective forces" can be understood as effective coherence gradients: tendencies of relational configurations to move toward states of higher global compatibility under constraints. In this paper we do not propose a gravitational model. We only note a structural resonance: gravitational dynamics in general relativity, unlike most forces in physics, is tightly tied to variational formulations, where realised histories and geometries are singled out by extremal principles under boundary conditions.

On this reading, one can speculatively view spacetime–matter configurations as relational candidates whose realised evolution is constrained optimisation of a functional (action-like, entropy-like, or information-like), in a way that is formally reminiscent of how probability fields *select* high-coherence configurations in other domains. Whether this analogy can be made precise—and whether it bears on quantum measurement rather than only on classical gravitational dynamics—remains open.

A more developed version of this idea would require: (i) a concrete coherence functional for spacetime configurations and its relation to known variational principles; (ii) a principled link, or a precise no-link, if any, between such a functional and probabilistic resolution in quantum measurement; (iii) empirical consequences in regimes where quantum and gravitational effects interact.

We leave these questions for future work.

# 6 Discussion

## 6.1 What is actually unified here?

At the structural level, we show that a range of systems can be described as relational probability fields shaped by context-dependent potentials and resolved by collapse or flow mechanisms. At an intermediate descriptive level, Section 2.4 introduces interface style descriptors (concentration, context sensitivity, and memory) that characterise how these fields respond to contextual change and how outcomes are realised. At a more speculative level, Section 5 proposes to interpret the associated probabilities as measures of relational coherence and to view collapse as a form of coherence optimisation.

The structural layer is the firm core: once we describe each theory in terms of fields of possibilities, relational potentials and selection rules, the formal similarities are straightforward. The coherence-optimisation layer is a hypothesis built on top of this structure. It suggests a way of reading the probabilities that appear in Born, Boltzmann, softmax, Bayes and replicator formalisms, but it does not change their empirical content.

Across all systems we considered—quantum measurement, Boltzmann ensembles, large language models, Bayesian inference and evolutionary game dynamics—the same minimal template appears:

1. A **field of possibilities**: a set of discrete alternatives (measurement outcomes, microstates, tokens, hypotheses, types).
2. A **relational potential**: a function that assigns each alternative a weight based on its relations to other variables and to a context (Hamiltonian and couplings, logits, likelihood, fitness, etc.).
3. **A mapping from potentials to an outcome distribution:** a rule that converts the induced structure into probabilities (Born, Boltzmann/softmax, Bayes, logit-like choice rules, etc.).
4. **A resolution mechanism (collapse or flow):** a process that realises particular outcomes or drives redistribution in time. Relational mode names the characteristic way these roles are coupled in a given domain, and it is operationalised through interface style descriptors (concentration, context sensitivity and memory).

In quantum mechanics, the Born rule connects a complex amplitude field to measurement statistics (Born, 1926; Landsman, 2009). In statistical mechanics, the Boltzmann distribution turns an energy landscape into equilibrium probabilities over configurations (Baxter, 1982). In large language models, the softmax function maps context-dependent scores to token probabilities (Bridle, 1990; Vaswani et al., 2017; Radford et al., 2019).

Bayesian inference reshapes a prior field over hypotheses via a likelihood that depends on data (Jaynes, 2003). Evolutionary game dynamics reweights the distribution over types via fitness functions that depend on population composition (Hofbauer & Sigmund, 1998, 2003).

The proposal of this paper is deliberately modest:

- We do **not** claim that these systems share the same ontology or microdynamics.
- We do **not** claim that quantum probability, statistical probability and algorithmic probability are identical.
- We do **not** derive the Born rule from the other examples, nor reduce quantum mechanics to classical or computational models.

What we claim instead is that once we describe these systems at the level of probability fields shaped by relational potentials, resolved by context-dependent selection, and characterised by interface style, they can be written in a shared language. In that language,

- the Born rule,
- the Boltzmann distribution,
- softmax token sampling,
- Bayesian updating, and
- replicator dynamics

all appear as different instances of the same core template (field, relational potential, mapping to probabilities, resolution/flow), with interface style providing a comparative layer and relational mode naming the characteristic way each domain instantiates this template.

This is a unification of form, not of substance. It does not replace the detailed theories in each domain. It provides a way to see them as instances of a more general pattern.

## 6.2 Relation to existing work

The structural framework is closely related to information-theoretic and statistical approaches that treat probability distributions as representations of constrained information states. The speculative coherence-optimisation hypothesis in Section 5 connects this to interpretations that emphasise relations, variational principles and informational or entropic origins of dynamics.

The idea that probability and information provide a unifying language across physics, statistics and computation is not new. Our proposal sits within this broader tradition, but makes a specific structural claim.

**Information-theoretic and statistical perspectives.**

Jaynes (2003) argued that probability theory is an extension of logic, and that many results in statistical mechanics can be seen as consequences of constrained maximisation of entropy. In that picture, probability distributions encode states of information, and the Boltzmann distribution emerges from constraints on expected energy. Our framework is compatible with this view, but emphasises the relational character of the underlying potentials: energies, likelihoods and scores are not intrinsic properties of isolated states, but functions of how these states relate to other variables and to a context. The interface-style descriptors introduced here are compatible with these traditions and can be seen as an additional comparative layer: they describe not the origin of a distribution, but how sharply it concentrates, how it reacts to contextual perturbations, and whether realised statistics exhibit history dependence.

**Relational interpretations of quantum mechanics.**

Relational quantum mechanics (Rovelli, 1996) suggests that the state of a system is defined only relative to another system, and that quantum events are about relations rather than observer-independent properties. Our use of the term "relational" is more limited and more formal: we describe how probabilities over outcomes are shaped by functions that depend on relations (e.g., couplings, payoffs, co-occurrence statistics). We do not commit to any particular interpretation of quantum mechanics; rather, we adopt a relational language because it captures a shared mathematical structure across domains.

**Derivations and reformulations of the Born rule.**

Many authors have attempted to derive the Born rule from other postulates of quantum theory, for example via decision-theoretic arguments (Deutsch–Wallace) or symmetry and "envariance" considerations (see Vaidman, 2019, for a review). Our approach does not belong to this family: we neither claim a derivation nor propose a modification of quantum postulates. Instead, we treat the Born rule as one example of a more general class of mappings from relational potentials to probability fields.

**Information, gravity and emergent forces.**

There is a growing body of work treating gravity and spacetime as emergent phenomena arising from informational or entropic principles (e.g., Verlinde's entropic gravity; Verlinde, 2011). More recently, Vopson (2019, 2022) has proposed a mass–energy–information equivalence principle, suggesting that information has a quantifiable mass and that information-theoretic considerations may underlie gravitational phenomena.

We do not build on these specific models and we remain agnostic about gravity in this paper. The connection is purely structural: these proposals also treat forces and macroscopic behaviour as emerging from the organisation of information and probability at a deeper level, which resonates with the coherence-based reading sketched in Section 5. Our contribution is to show that, even without committing to a specific physical mechanism, an abstract pattern of relational probability fields and collapse appears across several well-understood systems.

## 6.3 Scope and limitations

The framework proposed here has clear limitations, which we make explicit.

**Discrete fields and finite sets.**

For clarity, we formulated everything in terms of discrete sets of possibilities. Many relevant systems—quantum fields, continuous state spaces, neural dynamics—are naturally continuous or infinite-dimensional. Extending the framework to general measure spaces and continuous fields is conceptually straightforward but technically non-trivial, and we do not attempt it here.

**No claim of physical equivalence.**

The fact that different systems share a common formal pattern does not imply that they are physically equivalent or reducible to each other. Quantum amplitudes carry phase information and obey unitary dynamics; Boltzmann ensembles describe thermal equilibrium; large language models implement gradient-based optimisation on large text corpora. The framework abstracts away from these details and captures only the structure of how possibilities are weighted and selected.

**No new microphysical theory.**

We do not propose a new microscopic theory of quantum mechanics, gravity or neural computation. The framework is compatible with multiple interpretations of quantum mechanics and with existing physical models. It is a way of organising known mechanisms, not a replacement for them. Interface style is introduced as a comparative descriptor, not as an additional dynamical postulate.

**No direct claims about consciousness or phenomenology.**

Although similar field–collapse patterns appear in discussions of perception and decision-making, we deliberately avoid making claims about consciousness or subjective experience in this paper. Those questions require additional assumptions and empirical input that go beyond our present scope.

**Heuristic treatment of examples.**

Several of our examples (e.g., the GPT-2 token distribution for a particular prompt) are illustrative rather than exhaustive. A more complete study would require systematic empirical analysis of distributions and entropies across many prompts, temperatures, architectures and datasets, as well as more detailed statistical tests in Ising models and evolutionary games.

**The coherence-optimisation hypothesis is speculative.**

Section 5 proposed to interpret probabilities as relational coherence and to apply this idea across domains, including a tentative bridge toward gravitational dynamics. This goes beyond the purely structural claims of Sections 2–4 and should be read as a conjecture rather than as an established result.

These limitations are not defects but boundaries: they indicate that the present work should be read as a proposal for a common descriptive language, not as a finished theory.

## 6.4 Directions for further work

The structural perspective developed here opens several avenues for more detailed investigation:

**Formalising relational potentials.**

We used “relational potential” as an umbrella term for Hamiltonians, logits, likelihoods and fitness functions. One direction is to formulate a general mathematical notion of such potentials—e.g., as functions on graphs or hypergraphs of interactions—and to study what classes of probability fields and dynamics they can generate.

**Comparative entropy and information geometry.**

Since all our examples involve probability fields, tools from information theory and information geometry (e.g., entropies, divergences, Fisher metrics) could be used to compare interface style across systems (e.g., concentration/sharpness, context sensitivity, and memory), and how collapse or flow dynamics move distributions in these spaces.

**From coherence fields to effective forces.**

Explore whether coherence gradients can be formalised as effective dynamics in constrained relational configuration spaces, and whether gravitational variational principles can be expressed within that language.

**Conditions for specific mappings.**

In the examples studied here, different systems implement different mappings from potentials to probabilities (Born, Boltzmann, softmax, Bayes, replicator). It would be interesting to ask under what abstract conditions each of these mappings is singled out—for instance, whether certain symmetry or invariance requirements uniquely select the Born rule in a class of relational theories (Wallace, 2012; Vaidman, 2019).

**Interface style as an empirical/comparative layer.**

Identify measurable proxies for concentration, sensitivity and memory in each domain (e.g., temperature dependence and criticality in Ising systems; prompt perturbation tests in LLMs; decoherence-induced history dependence in quantum measurement models; path-dependence and hysteresis in non-equilibrium dynamics).

**Bridging to continuous and field-theoretic settings.**

Extending the framework to continuous configuration spaces and quantum field theory would test how far the "field–potential–collapse" pattern can be pushed beyond finite-dimensional toy models.

**Application to complex adaptive systems.**

Many real-world systems—neural networks in the brain, markets, ecosystems, multi-agent systems—can be viewed as distributions over configurations evolving under relational constraints. The present framework suggests that we might gain insight by explicitly identifying their probability fields, relational potentials and collapse/flow mechanisms and by asking whether they fit into the same structural family.

Our main claim is therefore not that there is a single "universal collapse law", but that a universal template exists for describing how structured fields of possibilities are related to realised outcomes. Making this template explicit may help to clarify connections between areas that are usually studied in isolation and to provide a common framework within which more specific, domain-dependent questions can be addressed.

# 7 Conclusion

We have argued that several superficially different systems—quantum measurements, Boltzmann ensembles, large language models, Bayesian inference and evolutionary game dynamics—can be described within a common structural framework. Each of them combines (i) a field of possibilities (outcomes, configurations, tokens, hypotheses, types), (ii) a relational potential that encodes how these possibilities are constrained or favoured in a given context, (iii) a mapping from relational potentials to an outcome distribution, and (iv) a resolution mechanism (collapse or flow) whose behaviour can be characterised by interface style—how strongly probability mass concentrates, how sensitively it responds to contextual perturbations, and whether effective outcome statistics exhibit memory. Taken together, these interface style descriptors operationalise a system's relational mode: its characteristic way of turning contextual relations into realised outcomes.

In quantum mechanics, the Born rule connects a complex amplitude field to measurement statistics. In statistical mechanics, the Boltzmann distribution turns an energy landscape into equilibrium probabilities. In large language models, a softmax function maps context-dependent logits to token probabilities. In Bayesian inference and evolutionary games, likelihoods and fitness functions play analogous roles in reshaping probability-like fields over hypotheses or types. At this level, our proposal is deliberately modest: we do not assert that these systems share a common substrate, nor that one can be reduced to another. Instead, we highlight a shared organisational pattern—structured fields of possibility, shaped by relations, mapped into distributions and resolved into realised outcomes—and suggest that many familiar probabilistic formalisms (Born, Boltzmann, softmax, Bayes, replicator) can be seen as instances of a single abstract template.

Built on top of this structural core, we sketched a tentative coherence-based interpretation. In that reading, the potentials that underlie these probability fields—energies, logits, likelihoods, fitness values and, in the quantum case, squared projection amplitudes—are not merely ad hoc weights but measures of relational coherence or fit between configurations and their context. Probabilities then quantify how strongly each configuration "belongs" in the relational structure defined by the system and its environment, while resolution can be viewed as a form of coherence optimisation on the space of possibilities. This does not eliminate stochasticity: individual outcomes remain random in each run, but the probability law governing that randomness is interpreted as encoding relational coherence rather than as an unexplained postulate. Importantly, the coherence-based reading does not by itself fix whether resolution is near-deterministic or highly stochastic; this depends on interface style (concentration and sensitivity) and on the specific realisation regime.

From this perspective, the Born rule becomes less mysterious. In Hilbert space, the natural measure of how strongly a state $|\psi\rangle$ relates to an eigenstate $|\lambda\_i\rangle$ is the squared inner product $|\langle\lambda\_i|\psi\rangle|^2$, singled out by the geometry of the space and by the requirement that probabilities over an orthonormal decomposition sum to one. Interpreting these Born weights as coherence scores aligns them conceptually with Boltzmann factors in energy landscapes and softmax scores in representation spaces: all three quantify a context-dependent notion of relational fit that shapes how probability mass is distributed over possibilities.

We also pointed out that this coherence-based view resonates with variational and information-theoretic approaches to gravity, in which spacetime geometries are characterised as extremising certain action or entropy functionals under relational constraints. If both quantum collapse and gravitational evolution can be cast, in different regimes, as forms of coherence optimisation on relational fields, then the gap between them may be narrower than it appears when they are treated as fundamentally separate theories. Instead of asking how to "add quantumness" to gravity or how to "add gravity" to quantum mechanics, one might ask whether both already implement a more general principle governing how relational structures resolve into realised configurations.

Many open questions remain. We have not provided a concrete coherence functional for quantum theory or for spacetime geometries, nor have we derived either the Born rule or Einstein's equations from a single optimisation principle. Extending the framework to continuous fields, specifying precise notions of relational coherence in different domains, and identifying empirical signatures—especially in regimes where quantum and gravitational effects interact—that could distinguish such a coherence-based framework from existing theories are all tasks for future work. Nevertheless, making the relational structure of probability fields explicit, and treating resolution as a candidate form of coherence optimisation, may help to shift the question from "why these probabilities?" to "what coherence principle do these probabilities express?" and to organise questions at the interface of quantum theory, statistical mechanics, machine learning and gravitational physics in a more unified way.

## References


Born, M. (1926). Zur Quantenmechanik der Stoßvorgänge. Zeitschrift für Physik, 37, 863–867. https://doi.org/10.1007/BF01397477

Baxter, R. J. (1982). Exactly solved models in statistical mechanics. Academic Press.

Breuer, H.-P., & Petruccione, F. (2002). The theory of open quantum systems. Oxford University Press.

Bridle, J. S. (1990). Training stochastic model recognition algorithms as networks can lead to maximum mutual information estimation of parameters. In D. S. Touretzky (Ed.), Advances in neural information processing systems (Vol. 2, pp. 211–217). Morgan Kaufmann.

Hofbauer, J., & Sigmund, K. (1998). Evolutionary games and population dynamics. Cambridge University Press.

Hofbauer, J., & Sigmund, K. (2003). Evolutionary game dynamics. Bulletin of the American Mathematical Society, 40(4), 479–519. https://doi.org/10.1090/S0273-0979-03-00988-1

Jaynes, E. T. (2003). Probability theory: The logic of science. Cambridge University Press.

Kullback, S., & Leibler, R. A. (1951). On information and sufficiency. The Annals of Mathematical Statistics, 22(1), 79–86.

Landsman, N. P. (2009). The Born rule and its interpretation. In D. Greenberger, K. Hentschel, & F. Weinert (Eds.), Compendium of quantum physics: Concepts, experiments, history and philosophy (pp. 64–70). Springer.

Neumaier, A. (2019). Born's rule and measurement. arXiv. http://arxiv.org/abs/1912.09906

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners (Technical report). OpenAI.

Rovelli, C. (1996). Relational quantum mechanics. International Journal of Theoretical Physics, 35(8), 1637–1678.

Shannon, C. E. (1948). A mathematical theory of communication. Bell System Technical Journal, 27(3), 379–423; 27(4), 623–656.

Vaidman, L. (2020). Derivations of the Born rule. In S. Gao (Ed.), Meaning of the Wave Function: In Search of the Ontology of Quantum Mechanics (pp. 283–305). Cambridge University Press.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In Advances in neural information processing systems (Vol. 30). Curran Associates.

Verlinde, E. (2011). On the origin of gravity and the laws of Newton. Journal of High Energy Physics, 2011(4), 29.

Vopson, M. H. (2019). The mass of information. AIP Advances, 9(9), 095206.

Vopson, M. M. (2022). Experimental protocol for testing the mass-energy-information equivalence principle. AIP Advances, 12(3), 035311.

Wallace, D. (2012). The emergent multiverse: Quantum theory according to the Everett interpretation. Oxford University Press.