

# From Field Dynamics to Extractive Systems: A Unified Account of Coarse-Graining, Cognition, and Economic Structure

Flyxion

*Independent Researcher*

April 19, 2026

## **Abstract**

This essay presents a unified interpretation of physical, cognitive, and economic systems through the lens of coarse-graining over an underlying scalar–vector–entropy field dynamics. Building on the Relativistic Scalar Vector Plenum (RSVP) framework and its TARTAN coarse-graining methodology, it is argued that probabilistic behavior, cognition, and platform economics are all manifestations of the same structural process: the projection of high-dimensional deterministic trajectories into constrained, observable forms. The key claim is not identity but shared invariant structure — these domains are governed by the same class of projection constraints and attractor dynamics, connected by structure-preserving functors rather than collapsed into a single ontological category. Extractive digital systems are then interpreted as pathological configurations of this projection, in which interface entropy at the interaction layer becomes a direct source of revenue. The analysis identifies three load-bearing structural claims: the selection of unistochastic rather than merely bistochastic transition matrices, the sheaf-like character of the coarse-graining projection, and the trajectory-level embedding of economic agents. Several problems are named explicitly as open.

## 1 Introduction

Modern theory treats physics, cognition, and economics as fundamentally distinct domains. Physics describes matter and energy, cognition describes mind and inference, and economics describes exchange and value. This separation, however, may be an artifact of representation rather than ontology.

The present work proposes that all three domains can be understood as instances of a single structural process: the compression of a high-dimensional dynamical system into a lower-dimensional observable interface, subject to consistency conditions across the compression boundary. Within this view, differences between domains arise from the scale and constraints of coarse-graining rather than from distinct underlying principles.

This claim requires precision. The framework does not assert a naive identity between physics, cognition, and economics. Rather, these are instances of systems governed by the same class of constraints and projection mechanisms, connected by structure-preserving functors. They share morphisms, invariants, and attractor structures while their ontological interpretations remain distinct. The unification is categorical and dynamical, not reductive.

Three structural claims are load-bearing for this argument. First, coarse-graining over RSVP trajectories selects unistochastic rather than merely bistochastic transition operators, and this selection depends on the retention of coherent phase structure under projection. Second, the coarse-graining map is sheaf-like rather than a simple averaging operation, and this is what allows it to support non-Markovian dynamics and global consistency conditions. Third, economic agents are best embedded at the level of trajectories through coarse-grained state space, not as static regions or instantaneous states, and this determines how the extraction problem is formulated.

Several further questions — the precise form of higher cohomological obstructions, the detailed construction of the CLIO manifold, the variational formulation of extraction, and the status of  $\hbar$  as an emergent scale — are genuinely open and are named as such rather than papered over.

## 2 RSVP as a Generative Substrate

The Relativistic Scalar Vector Plenum (RSVP) framework models reality as a coupled field system. The full microstate of the system is given by

$$X(t) = (\Phi(x, t), \vec{v}(x, t), S(x, t)),$$

where  $\Phi : \mathcal{M} \rightarrow \mathbb{R}$  is a scalar potential encoding energy density,  $\vec{v} : \mathcal{M} \rightarrow T\mathcal{M}$  is a vector field representing directional flow, and  $S : \mathcal{M} \rightarrow \mathbb{R}_{\geq 0}$  is an entropy density.

The base manifold  $\mathcal{M}$  is signature-flexible. In regimes requiring relativistic interpretation it is taken to be Lorentzian; in purely computational or cognitive reductions a Riemannian or abstract measure-theoretic base suffices. Compactness is not assumed globally; one works with compact subsets with boundary conditions induced by the coarse-graining scale.

The three fields are not independent. They form a constrained triple coupled dynamically through transport and entropy production equations. The entropy density  $S$  is partially determined by gradients and torsion in  $\vec{v}$  and by curvature in  $\Phi$ , but no single algebraic equation of state is imposed universally. The constraint is variational and dynamical rather than pointwise. Non-negativity of  $S$  is not imposed arbitrarily; it follows from its role as a coarse-grained measure of microstate multiplicity.

The dynamics evolve the triple  $(X(t))$  deterministically according to coupled partial differential equations. Observable structure does not arise directly from  $X(t)$  but from projections of it. These projections are constrained by thermodynamic consistency and limited access to information.

### 3 TARTAN Coarse-Graining and the Sheaf Structure of Projection

The TARTAN (Trajectory-Aware Recursive Tiling with Annotated Noise) framework formalizes the projection through recursive tiling and memory-bearing perturbations. Let  $\{T_i\}$  be a covering of  $\mathcal{M}$  by tiles. This covering is not a strict partition: overlaps are permitted and are necessary to encode consistency conditions across scales. At finer recursion levels tiles behave approximately as partitions; at coarser levels the overlaps carry the memory necessary for non-Markovian dynamics.

The coarse-grained variable over tile  $T_i$  is

$$\chi_i(t) = \mathcal{F}(X|_{\Omega_i}, \eta_i),$$

where  $\Omega_i$  denotes the region associated with  $T_i$  and  $\eta_i$  is annotated noise. Crucially,  $\eta_i$  is not independent stochastic input. It encodes a memory kernel inherited from prior tilings, preserving structured remnants of past states rather than introducing exogenous randomness. The memory kernel is generally non-integrable, though it may admit effective decay in certain regimes.

The map  $\mathcal{F}$  is not assumed linear. It is measurable and typically nonlinear, reflecting aggregation, thresholding, and constraint-based projection.

The critical structural feature of TARTAN is that the projection  $\pi : X(t) \rightarrow \chi_i(t)$  is sheaf-like rather than a simple marginalization. Overlapping tiles enforce partial agreement conditions on their shared boundaries, functioning as local gluing constraints. This has two consequences. First, the projection is inherently non-Markovian, because the compatibility conditions implicitly encode history: intermediate states do not form a sufficient statistic for future evolution. Second, the projection retains enough structure to support coherent phase transport, which is what selects unistochastic rather than merely bistochastic dynamics at the coarse-grained level.

The non-Markovian character follows jointly from the memory encoded in  $\eta_i$  and from geometric features of  $\vec{v}$ , specifically from torsion  $\vec{T} = \nabla \times \vec{v}$ . Neither alone is sufficient in general; it is their interaction that produces persistent temporal correlations. Non-divisibility of the transition structure then follows from the coupling between entropy gradients and vector torsion: divisibility would require intermediate states to be sufficient statistics, which they are not when torsion is present.

#### 4 Phase Space, Path Integrals, and the Projection Boundary

The projection  $\pi : X(t) \rightarrow \chi_i(t)$  can be understood more precisely by lifting the RSVP dynamics into a path space formulation. Rather than treating  $X(t)$  as a trajectory in configuration space alone, we consider the ensemble of all admissible paths  $\gamma$  consistent with the field dynamics.

Define the path functional

$$\Psi[\gamma] = \exp\left(-\int_{\gamma} \mathcal{L}(X, \partial X) d\tau\right),$$

where  $\mathcal{L}$  is an effective Lagrangian incorporating entropy gradients, vector torsion, and scalar curvature contributions.

The coarse-graining map then acts not on individual states but on equivalence classes of paths. The projection boundary is therefore not a spatial boundary alone but a boundary in path space: a restriction on which trajectories are distinguishable under  $\mathcal{F}$ .

This clarifies the emergence of probability. Observable transition frequencies arise from integrating  $\Psi[\gamma]$  over equivalence classes of paths, while phase coherence arises from the composition properties of these path weights. The distinction between classical stochasticity and quantum-like behavior is therefore a property of how path equivalence is defined under projection, not a property of the underlying dynamics itself.

## 5 On the Emergence of Inner Product Structure

The appearance of Hilbert space structure in certain regimes is not assumed but emerges from symmetry conditions on the projection.

An inner product on the space of coarse-grained amplitudes becomes well-defined when two conditions hold simultaneously. First, the coarse-graining preserves phase coherence across overlapping tiles, so that amplitudes can be consistently compared across boundaries. Second, the projection admits a symmetry under reversal or recombination of admissible paths, ensuring that transition amplitudes form a closed algebra under composition.

Under these conditions one may define

$$\langle \psi, \phi \rangle = \sum_i \overline{\psi_i} \phi_i,$$

where  $\psi_i$  and  $\phi_i$  are amplitude assignments to coarse-grained states. This structure is not fundamental but induced by the projection geometry: it is a consequence of phase coherence and compositional symmetry, not a prior assumption.

Outside such regimes, the appropriate structure is more general: a sheaf of state spaces  $\mathcal{H}$  without a globally defined inner product. Hilbert space is therefore a special case corresponding to a high degree of symmetry and coherence in the underlying projection. The regimes in which it emerges are precisely those in which the unistochastic description of Section 5 is most reliable.

The emergence of quantum-like transition probabilities from a classical field substrate requires a careful distinction. A purely entropy-preserving coarse-graining would produce a bistochastic transition matrix — one that preserves the uniform distribution — but not necessarily a unistochastic one. Every unistochastic matrix is bistochastic, but the converse fails. Selecting the unistochastic subclass requires an additional structural ingredient.

In RSVP, this ingredient is the circulation structure of the vector field. Torsion and path-dependence in  $\vec{v}$  mean that coarse-graining does not merely count transition frequencies. It retains residual orientation information from line integrals of  $\vec{v}$  along admissible paths. This induces a phase functional  $\Omega[\gamma]$  over trajectories.

**Definition 1** (Phase functional). For an admissible path  $\gamma$  in the microstate space, the phase functional is

$$\Omega[\gamma] = \oint_{\gamma} \vec{v} \cdot d\ell,$$

where the integral is taken along the trajectory projected onto  $\mathcal{M}$ .

When the projection respects composition of paths, these phases compose additively. The effective transition amplitudes therefore define a representation of the path composition semi-

group on amplitude space, and since the projection preserves total probability across all admissible compositions, this representation must be norm-preserving, forcing it into the unitary class rather than the broader class of contractive or isometric operators. At that point the framework is no longer working with probabilities but with amplitudes whose squared moduli give probabilities, and the consistency of composition forces the amplitude operator to be unitary rather than merely stochastic.

**Proposition 1** (Unistochasticity from phase-coherent projection). *If the coarse-graining map  $\pi$  preserves the composition structure of paths and retains phase information via  $\Omega[\gamma]$ , then the effective transition operator  $U$  acting on amplitude space is unitary, and the coarse-grained transition probabilities satisfy*

$$P_{ij} = |U_{ij}|^2.$$

The essential point is that unistochasticity is selected because the coarse-graining is not purely measure-theoretic. It is measure plus coherent transport structure. Without the retained phase, the projection collapses back to bistochasticity. This is where the RSVP framework recovers quantum structure rather than merely classical stochasticity: the vector field carries phase information that survives projection in a constrained and composition-consistent way.

The precise conditions under which convergence to  $|U_{ij}|^2$  holds, and whether the unitary  $U$  here is identical to the one arising from the Madelung complexification  $\Psi = \sqrt{\rho} e^{i\Omega/\hbar}$ , remain open structural problems. The two constructions are expected to be related by the projection framework, but the identification has not been established in full generality.

## 6 Memory, Decay, and Effective Markovianity

The memory kernel encoded in  $\eta_i$  is in general non-integrable, reflecting the persistence of historical information in the coarse-grained projection. However, in certain regimes the influence of past configurations decays sufficiently rapidly that the dynamics admit an effective finite-memory approximation.

Formally, there exists a timescale  $\tau$  such that for  $t \gg \tau$ ,

$$P(\chi_{t+1} | \chi_t, \chi_{t-1}, \dots) \approx P(\chi_{t+1} | \chi_t, \dots, \chi_{t-k})$$

for some finite  $k$  determined by the coarse-graining depth. This does not restore true Markovianity but defines an effective Markov order at the scale of observation. The transition from non-Markovian to effectively Markovian behavior is therefore a function of the projection regime rather than a property of the underlying dynamics.

This observation carries consequences in both physics and economics. In physical systems it explains why classical stochastic models can approximate quantum systems in certain lim-

its without capturing their full structure. In economic systems it explains why short-term interactions can appear memoryless even when long-term trajectory dependence is structurally dominant — and why intervention at the short-term level alone is insufficient to alter extraction dynamics whose roots lie in trajectory history.

The same sheaf-like projection structure appears in cognition, and the identification is not merely analogical. Cognition, under this framework, is the process of extending locally consistent projections into a globally consistent section.

Local cognitive data is identified with the tile variables  $\chi_i$ , which are coarse-grained representations of  $(\Phi, \vec{v}, S)$  over regions  $\Omega_i$ . In higher-level abstractions these may be replaced by symbolic or semantic objects derived from them, but their mathematical role remains the same: they are local sections of a sheaf of state spaces  $\mathcal{H}$  over  $\mathcal{M}$ .

The sheaf  $\mathcal{H}$  is a sheaf of local state spaces — the spaces in which coarse-grained configurations live. It is distinct from the sheaf of observables  $\mathcal{O}$ , which is a derived or dual sheaf whose sections are measurement functionals acting on states. The obstruction class measuring gluing failure lives in the cohomology of  $\mathcal{O}$ , not of  $\mathcal{H}$  directly: one asks whether locally consistent observation records can be assembled into a globally consistent observable, not merely whether state-space elements agree on overlaps.

The sheaf  $\mathcal{H}$  is understood as a sheaf of state spaces rather than strictly of Hilbert spaces. Hilbert structure emerges in regimes where phase coherence is preserved under projection and inner products become well-defined, as characterised in Section 6. The Spherepop construction, in which local state transitions correspond to amplitude relaxations on a five-dimensional lattice, represents a special case in which this Hilbert structure is made explicit.

**Proposition 2** (Gluing condition). *A consistent global cognitive state exists if and only if the local sections  $\{\chi_i\}$  agree on all overlaps  $\Omega_i \cap \Omega_j$ . Failure of this condition is measured by a cohomological obstruction.*

The first obstruction class lies in  $H^1(\mathcal{M}, \mathcal{O})$ , where  $\mathcal{O}$  is the sheaf of local observables. This class captures pairwise gluing failures: cases where two locally consistent descriptions cannot be reconciled on their shared boundary.

However,  $H^1$  does not exhaust the possible obstructions. Multi-scale inconsistency — where local pairs are compatible but triple overlaps fail — may require  $H^2$  or higher. Full classification of cognitive incoherence in this framework likely requires derived category methods. This is named here as an open problem rather than a settled result.

The practical consequence is that what one calls hallucination, unresolved ambiguity, or inconsistent belief corresponds formally to the non-existence of a global section. The cognitive system is attempting to solve a sheaf extension problem, and the obstructions are the invariants that measure how badly it fails.

CLIO (Cognitive Loop via In-Situ Optimization [1]) acts on this process as an adaptive update operator. In quantum-like regimes it admits representation as a completely positive trace-preserving (CPTP) map, but more primitive formulations exist in terms of constraint propagation over tile boundaries. CLIO does not operate on  $\mathcal{M}$  directly but on a derived manifold of representations constructed over trajectory space. Fixed points of CLIO correspond to globally coherent states; multiple locally stable fixed points may coexist, separated by regions of high  $S$  or  $|\nabla\Phi|$  or by topological obstructions. Which of these mechanisms governs transitions between fixed points is an open question.

## 7 Topological and Energetic Separation of Fixed Points

The coexistence of multiple CLIO fixed points raises the question of how they are separated and how transitions between them occur. Two mechanisms are possible and should be distinguished.

In energetic separation, fixed points are divided by regions of high entropy density  $S$  or large gradients  $|\nabla\Phi|$ . Transitions between them are thermally or dynamically activated: a sufficient fluctuation in the underlying field dynamics can drive the system across the barrier. Such transitions are unlikely but continuous in the sense that no topological obstruction prevents them.

In topological separation, fixed points lie in distinct homotopy classes of sections of  $\mathcal{H}$ . No continuous deformation connects them without violating local consistency conditions on tile overlaps. Transitions between such states are not merely unlikely but structurally obstructed: they require a discontinuous reorganization of the global section, which in cognitive terms corresponds to a qualitative shift in the structure of understanding rather than a gradual update.

In general both mechanisms may be present simultaneously. The distinction matters for how one models cognitive change: energetically separated fixed points admit gradual revision, while topologically separated ones require something closer to a gestalt shift. Determining which mechanism dominates in a given regime, and whether the two can be distinguished observationally, is an open problem.

The embedding of economic systems into this framework requires a precise decision about the level of abstraction at which agents are represented. The correct level is neither the raw tile region  $\Omega_i$  nor the instantaneous variable  $\chi_i$ , but the trajectory through tile space.

**Definition 2** (Agent as trajectory). An agent is an equivalence class of paths  $[\gamma]$  in the coarse-grained state space, where

$$\gamma = (\chi_{i_0}(t_0), \chi_{i_1}(t_1), \dots)$$

with successive states connected by admissible transitions in the TARTAN dynamics, and two paths are equivalent if they are related by an admissible refinement or reparameterization of the tiling. This equivalence removes dependence on the specific resolution at which the tiling is defined, consistent with the scale-relative stance of the framework.

This definition prevents identity from being treated as a static container. It is the trajectory-level structure that gives identity its economic meaning: the history of commitments, fulfillments, and outcomes that constitutes a participant’s record. Economic actions are then morphisms between segments of such trajectories, and markets are interaction patterns between paths rather than between nodes.

The platform layer constitutes a secondary projection that maps these trajectories into visibility, ranking, and matching structures. In a productive regime this projection preserves the trajectory structure that makes matching meaningful. In an extractive regime the projection is distorted: revenue becomes proportional not to successful path completion but to fragmentation of trajectories — to broken, delayed, or mutually incompatible paths.

This connects directly to the entropy argument. Interface entropy  $H(t)$ , defined over the distribution of active proposals at time  $t$ , is not identical to the thermodynamic entropy  $S$  but is an induced quantity derived from it under projection into the interaction space. The two are related by the mapping from field configurations to distributions over observable actions.

**Condition 1** (Non-extraction from entropy). *A non-extractive platform satisfies*

$$\frac{\partial R}{\partial H} \leq 0,$$

*where  $R$  is platform revenue and  $H$  is interface entropy. Infrastructure revenue is non-increasing in interface disorder.*

Under this condition the platform cannot profit from failed or fragmented interactions. Its incentive structure aligns with path completion rather than path fragmentation.

The bond conservation condition formalizes this at the level of individual proposals: unresolved interactions return their stake to participants rather than converting it to platform revenue. Together these conditions define a regime outside the extraction attractor, where the locally optimal strategy for infrastructure is to reduce mismatch rather than amplify it.

These conditions are currently imposed as design constraints. A variational formulation is expected to exist, in which extraction corresponds to a non-conservative flow violating an entropy-consistent action principle defined over the space of projection operators rather than over field configurations directly. This is an open problem.

## 8 The Extraction Attractor as a Pathological Projection

The convergence of diverse systems toward extractive behavior is interpretable within this framework as the emergence of a shared attractor in the space of coarse-graining mechanisms.

When a platform accumulates scale, two changes occur simultaneously. The cost of exit for participants increases as their trajectory histories become bound to the platform’s internal projection. And control over the secondary projection — the mapping from trajectories to visibility — becomes concentrated in a single authority.

Once both conditions hold, the platform’s incentive structure undergoes a phase transition. Below the threshold, the platform competes for participation by producing value. Above it, the platform manages the dependency already created. The transition is largely irreversible because the accumulation of trajectory histories is itself what makes the high-extraction equilibrium stable.

In RSVP terms, this corresponds to a regime in which entropy at the interface layer becomes a positive contributor to revenue. The gradient condition  $\partial R/\partial H \leq 0$  is violated. The platform begins to profit from the fragmentation of trajectories, and the sheaf-like consistency conditions that would support global coordination are no longer maintained.

The attractor is stable not because any actor designed it but because the competitive environment selects for it. Platforms that do not evolve toward extraction produce positive externalities that competitors exploit without contributing. They cannot recapture the value generated by supporting trajectory completion. Over time the space is populated by platforms that have adopted extraction strategies, either by design or evolutionary pressure.

Cognitive infrastructure — AI systems embedded in workflows — is entering the same trajectory. The chokepoint shifts from access to work (labor platforms) or access to attention (social networks) to access to the cognitive projection itself: the capacity to perform structured inference over a semantic field. If that projection becomes mediated and priced, the extraction attractor extends into the layer where trajectories of meaning are formed.

## 9 Toward Non-Extractive Architectures

If extraction is a property of how projections are structured, alternatives must be designed at the same level. The conditions derived above suggest the outlines of such an architecture.

Proposals should be modeled as typed commitments rather than purchased visibility — morphisms between trajectory segments, not access tokens. The platform cannot own the act of matching; it observes, indexes, and interprets a shared event substrate that it does not constitute.

Identity must be bound to trajectory rather than to namespace. The economic value of an identity is not separable from the history that constitutes it. A name without the trajectory is inert; leasing the namespace does not transfer the trajectory.

The matching layer should be decomposable: a plurality of interpretation functors  $F_i : \mathcal{E} \rightarrow \mathcal{R}_i$  operating over a shared event category  $\mathcal{E}$ , rather than a single centralized ranking function. Control over relevance is distributed across interpreters rather than concentrated in a single projection authority.

Bond conservation ensures that infrastructure cannot profit from unresolved proposals. The non-extraction condition ensures that amplifying interface entropy does not increase revenue. Together they define a projection regime in which the extraction attractor is no longer locally optimal, because the mechanisms through which extraction operates are structurally unavailable.

Whether this regime can be sustained at scale against the competitive pressure of extractive alternatives is the central open empirical question. The architecture is necessary but not sufficient; the social and competitive conditions that would allow it to survive before it scales must also be present.

## 10 Functorial Structure and Multi-Domain Correspondence

The unification proposed in this essay is best formalized in categorical terms. Let  $\mathcal{C}_{\text{phys}}$ ,  $\mathcal{C}_{\text{cog}}$ , and  $\mathcal{C}_{\text{econ}}$  denote categories corresponding to physical, cognitive, and economic systems respectively. Objects in each category are coarse-grained state spaces at the appropriate level; morphisms are admissible transitions or projections preserving the relevant invariants.

These categories are not identical but are connected by structure-preserving functors

$$F_{\text{pc}} : \mathcal{C}_{\text{phys}} \rightarrow \mathcal{C}_{\text{cog}}, \quad F_{\text{ce}} : \mathcal{C}_{\text{cog}} \rightarrow \mathcal{C}_{\text{econ}}.$$

Each functor preserves projection structure, non-Markovian dependence, and attractor behavior. What changes across domains is the interpretation of objects and morphisms, not the underlying form of the dynamics or the class of constraints that govern them.

This formulation makes the scope of the unification claim precise. It avoids collapsing all domains into a single category while retaining a rigorous sense in which they are governed by the same structural laws. The equivalence is in the invariants and attractors, not in the ontological identity of the objects.

The analysis has argued that physics, cognition, and economic systems share invariant structure at the level of their projection mechanisms. All three involve the compression of high-dimensional dynamical information into lower-dimensional observable interfaces, subject to

consistency conditions across the compression boundary. All three exhibit attractor dynamics that depend on whether those conditions are maintained or violated.

The three load-bearing structural claims have been made explicit. Unistochasticity is selected by the retention of coherent phase transport under projection, not by entropy preservation alone. The sheaf-like character of the projection is what enables non-Markovian dynamics and global consistency conditions. Agents are trajectories through coarse-grained state space, and extraction is the distortion of secondary projections that makes revenue proportional to trajectory fragmentation rather than trajectory completion. These three claims correspond respectively to structure at three levels: microdynamic (phase transport selecting unistochasticity), mesoscopic (sheaf projection enabling constraint closure and cognition), and macroscopic (trajectory embedding determining the form of economic extraction).

Several problems remain open: the identification between the two constructions of the unitary  $U$ , the full cohomological classification of cognitive incoherence, the variational formulation of the extraction condition, the detailed construction of the CLIO manifold, and the stabilization of  $\hbar$  as an emergent scale across different projection regimes. These are not gaps to be concealed but directions from which the framework can be extended.

The problem of platform economics, in this light, is not separate from the problem of quantum foundations or cognitive coherence. All are instances of the same question: under what conditions does structure survive compression, and what is destroyed when it does not?

## References

- [1] N. Cheng, G. Broadbent, and W. Chappell. Cognitive Loop via In-Situ Optimization: Self-Adaptive Reasoning for Science. arXiv:2508.02789, 2025.
- [2] J. A. Barandes. *The Stochastic-Quantum Correspondence*. arXiv:2302.10778, 2023.
- [3] J. A. Barandes. *New Prospects for a Causally Local Formulation of Quantum Theory*. arXiv:2402.16935, 2024.
- [4] J. Barbour. *The End of Time*. Oxford University Press, 1999.
- [5] D. Bohm. A suggested interpretation of the quantum theory in terms of “hidden” variables. *Physical Review*, 85(2):166–193, 1952.
- [6] R. Zwanzig. *Nonequilibrium Statistical Mechanics*. Oxford University Press, 2001.