THE ENTROPY OF CHOICE

Mathematical Appendices to the Boundary-Complexity Hypothesis

Rajesh Singh
MetaForgeX AI (Pty) Ltd | NSDM Institute
Johannesburg, South Africa
Working Paper 2026-02A | Revision 2.0 | July 2026

RESEARCH STATUS
Independent mathematical supplement. Not yet peer-reviewed. Appendix F contains the strongest result: a covering-number upper bound under explicit regularity, metric and local-complexity assumptions. Appendices G-I present structured correspondences, not proven equivalences. Appendices J-L are research questions rather than established quantum or hardware predictions. Appendices M-N define operational hypotheses and falsification tests.

AI-assistance disclosure
AI was used for research support, drafting and adversarial review. The author directed the argument, selected the claims, reviewed the sources and accepts responsibility for the final text.

Notation and scope

Let M be a compact metric state-action manifold, d its ambient dimension, and pi a binary policy. The action-changing set is denoted by partial D_pi. Its intrinsic dimension is k, which need not equal d-1 in irregular or lower-dimensional cases. H^k denotes k-dimensional Hausdorff measure. Epsilon is a declared observational or representational resolution. N(epsilon, Pi_delta, d_H) is the covering number of a near-optimal policy class Pi_delta under Hausdorff boundary distance d_H. All logarithms are natural unless stated otherwise.

These appendices distinguish four levels of claim:

- Derived result: follows from stated assumptions.
- Restricted correspondence: follows only for a special model family or approximation.
- Structural analogy: shares a mathematical form but not a common physical mechanism.
- Empirical hypothesis: must be tested and may fail.

Appendix F. Boundary covering bound

This appendix states the mathematically strongest version of the proposal. It does not assume that every patch of a decision boundary varies independently, and it does not fix the boundary dimension at two.

F.1 Near-optimal policy class

Let X be distributed according to P_X on M, let pi* be an optimal policy, and let L be a decision loss. Define

$$\text{Pi\_delta} = \{ \text{pi} : \text{E\_X}[\text{L}(\text{pi}(\text{X}), \text{pi*}(\text{X}))] <= \text{delta} \}.$$

The class Pi_delta contains policies whose expected loss relative to pi* is at most delta. The bound concerns distinguishable boundaries inside this class, not arbitrary geometric surfaces.

F.2 Regular boundary class

Assume that every partial D_pi belongs to a regular k-dimensional class B with uniformly bounded local configuration complexity. Informally, a boundary can be covered by epsilon-scale patches, and each patch admits at most Lambda distinguishable local states once smoothness, curvature or other regularity constraints are imposed.

F.3 Proposition

PROPOSITION F.1 - Boundary covering bound

If the boundaries in Pi_delta have uniformly bounded k-dimensional Hausdorff measure and local epsilon-scale configuration count Lambda, then the number of pairwise epsilon-distinguishable policy boundaries satisfies the bound below for a constant C determined by the boundary class and metric.

log N(epsilon, Pi_delta, d_H) <= C * sup_{pi in Pi_delta} H^k(partial D_pi) / epsilon^k.

Proof sketch. Cover each boundary by at most C_0 H^k(partial D_pi)/epsilon^k metric balls. If each local patch admits at most Lambda admissible states, then the total number of configurations is at most Lambda raised to the covering count. Taking logarithms yields the result with C = C_0 log Lambda. Correlation between patches can only reduce the actual configuration count, so the expression remains an upper bound within the assumed class.

F.4 What the proposition does and does not say

- It gives a geometric upper bound on distinguishable policy boundaries at a declared resolution.
- It does not prove that every model requires that much capacity.
- It does not establish a universal parameter-count formula.
- It depends on the metric and representation used to measure the boundary.
- The epsilon^-2 expression is only the special case k = 2.

Appendix G. Rate-distortion localisation near a decision boundary

The original appendices described an exact equivalence with Shannon rate-distortion theory. The defensible claim is narrower: under a local value-gap model, decision regret is concentrated in a shell around the action-changing boundary, producing a related boundary-sensitive scaling.

G.1 Decision distortion

D(pi_hat) = E_X[L(pi_hat(X), pi*(X))].

A rate-distortion problem asks for the smallest information rate needed to keep expected decision loss below D. The relevant distortion is task loss or regret, not geometric displacement by itself.

G.2 Boundary-shell approximation

Suppose the value gap between competing actions grows approximately linearly with normal distance r from a regular boundary: Delta V(x) approximately kappa |r|. A perturbation changes the action only when it crosses the boundary. For small perturbations, material regret is therefore concentrated in a tubular neighbourhood of partial D.

P_X(error with action change) approximately integral over the boundary shell.

If P_X has a smooth density and the shell thickness is eta, its probability mass is approximately eta times an integral of the density over the boundary. This creates a principled link between distortion and boundary-local mass.

G.3 Status of the correspondence

RESTRICTED CORRESPONDENCE
The result supports boundary localisation of decision distortion. It does not establish R(D) = S_choice + O(1) for general sources, losses or policy classes. A full equivalence would require solving the rate-distortion problem for a specified source distribution, policy family and distortion function.

Appendix H. Boundary-localised variational model

Variational free energy provides a tractable upper bound on surprise in Friston's formulation. The present appendix does not claim a general free-energy/area theorem. It examines a restricted posterior family concentrated near a regular decision boundary.

H.1 Tubular posterior family

$$q(\theta) = Z^{-1} \exp\left(-\rho(\theta)^2 / (2 \sigma^2) \right),$$

where rho is signed distance to partial D and sigma is the normal width. Under a regular tube approximation,

$$Z \approx H^k(\partial D) * \sigma * \sqrt{2 \pi},$$

up to curvature and density corrections. The differential entropy therefore contains a term log H^k(partial D) plus normal-width and constant terms.

H.2 Correct interpretation

- For this posterior family, the entropy term depends on boundary measure.
- The energy term remains model-dependent.
- This does not prove a global upper bound on variational free energy by boundary area.
- Precision changes the normal width sigma, but it does not establish universal cultural or psychological differences.

Appendix I. Holographic and Bekenstein-Hawking correspondence

Black-hole entropy is a physical result in gravitational thermodynamics: at leading semiclassical order it is proportional to horizon area. The proposed decision bound is a metric-entropy statement about a policy class. They do not arise from the same physical mechanism.

Feature	Black-hole setting	Decision-boundary setting
Boundary object	Event horizon	Action-changing set in a chosen representation
Measure	Physical horizon area	Hausdorff measure under a task metric
Resolution	Planck-scale physics	Declared observational or representational epsilon
Entropy meaning	Thermodynamic / quantum-gravitational state count	Log covering number of a policy class
Status	Established physical theory within its domain	Proposed mathematical model requiring empirical validation

The safe conclusion is a shared area-over-resolution form. The paper should not state that the two bounds are the same theorem, that gravity has a decision-theoretic analogue, or that black-hole saturation transfers to machine-learning policies.

Appendix J. Quantum-kernel research question

Quantum feature maps can embed inputs into high-dimensional Hilbert spaces. This representational fact does not establish a computational advantage. A valid quantum-advantage claim requires a task, an efficient encoding, a quantum algorithm, a classical lower bound, measurement cost and a noise model.

CONDITIONAL HYPOTHESIS J.1
Boundary complexity may help identify decision tasks worth testing with quantum kernels, but no qubit threshold, date or advantage follows from the covering bound alone.

- Estimate boundary complexity with a representation-stable metric.
- Compare against strong classical kernels and neural baselines.
- Include data-loading and measurement cost.
- Require reproducible end-to-end advantage, not feature-space dimension alone.

Appendix K. Thermodynamic-computing research question

Thermodynamic and probabilistic hardware may be well suited to sampling or energy-based inference. The proposed decision-boundary theory does not currently derive a hardware-efficiency law. Efficiency must be measured in joules, latency, throughput, solution quality and total system cost on a declared workload.

CONDITIONAL HYPOTHESIS K.1
If a decision workload is dominated by sampling near competing actions, probabilistic hardware may offer an efficiency advantage. The scaling of that advantage with boundary complexity is an open empirical question.

Appendix L. No qubit-count timeline theorem

The original timeline extrapolated a doubly exponential trend and mapped boundary configuration count directly to qubit count. That inference is not valid. Logical qubits, circuit depth, error correction, state preparation and algorithmic complexity cannot be collapsed into Hilbert-space dimension.

REVISION
Remove the 2028-2030 prediction. Replace it with a technology-watch protocol that records hardware capabilities and triggers a benchmark only when a specific algorithm can be executed against a competitive classical baseline.

Appendix M. Model-capacity knee hypothesis

The original Model-Selection Theorem assumed that representational capacity maps directly to parameter count and that any model above the covering number achieves minimum regret. That is not established. Optimisation, architecture, inductive bias, data coverage and training procedure all matter.

M.1 Operational hypothesis

HYPOTHESIS M.1
For a controlled model family trained on the same data and objective, performance may exhibit a capacity knee. The location of that knee may correlate with an independently estimated measure of boundary complexity.

M.2 Valid test design

- Choose model families with ordered, measured capacity.
- Estimate boundary complexity before inspecting the performance curve.
- Fit a segmented or saturating curve with uncertainty intervals.
- Compare prediction against input dimension, sample size, margin and simpler complexity measures.
- Repeat across seeds and unrelated domains.
- Treat failure to outperform these baselines as evidence against the operational theory.

Appendix N. Empirical falsification protocol

Test	Required evidence	Failure condition
N.1 Synthetic recovery	Known boundaries are correctly ordered by the estimator; random-label controls fail.	Estimator cannot recover known ordering or responds strongly to random labels.
N.2 Boundary-shell localisation	Material errors or regret concentrate nearer the true action boundary.	No localisation after controlling for class frequency and density.
N.3 Boundary vs irrelevant dimension	Capacity tracks boundary complexity more strongly than added nuisance dimensions.	Nuisance dimension predicts capacity equally well or better.
N.4 Representation stability	Estimator remains stable under declared bi-Lipschitz or information-preserving	Estimate changes arbitrarily under

	transforms.	equivalent representations.
N.5 Capacity-knee prediction	Pre-registered boundary estimate predicts the smallest near-best model.	Prediction misses across at least three unrelated domains.
N.6 External replication	Independent team reproduces direction and uncertainty of the main result.	Result depends on one dataset, seed or estimator choice.

N.1 Proposed sequence

- Stage 1: synthetic geometric wind tunnel.
- Stage 2: the existing QSR decision-state benchmark.
- Stage 3: public uplift, predictive-maintenance, routing and decision datasets.
- Stage 4: preregistered cross-domain replication.
- Stage 5: only then evaluate hardware-routing implications.

Closing limitations register

Limitation	Why it matters	Research response
Representation dependence	Boundary measure can change under latent reparameterisation.	Define and test a stable metric class.
Boundary observability	Labels may not identify the true action-switching set.	Use synthetic ground truth and causal/decision outcomes.
Patch dependence	Local boundary states are correlated.	Use regularity classes and empirical metric entropy.
Distribution dependence	Low-density boundary regions may matter little operationally.	Weight geometry by P_X and regret.
Optimisation gap	A model may have capacity but fail to learn the boundary.	Separate expressivity, optimisation and data errors.
Multi-action decisions	Boundaries can be stratified and intersect.	Extend to multiclass boundary complexes.
Dynamic decisions	The boundary can drift over time.	Add temporal geometry and monitoring.

References

Bekenstein, J. D. (1973). Black holes and entropy. Physical Review D, 7(8), 2333-2346. https://doi.org/10.1103/PhysRevD.7.2333

Bousso, R. (2002). The holographic principle. Reviews of Modern Physics, 74(3), 825-874. https://doi.org/10.1103/RevModPhys.74.825

Friston, K. (2010). The free-energy principle: a unified brain theory? Nature Reviews Neuroscience, 11, 127-138. https://doi.org/10.1038/nrn2787

Schuld, M., & Killoran, N. (2019). Quantum machine learning in feature Hilbert spaces. Physical Review Letters, 122, 040504. https://doi.org/10.1103/PhysRevLett.122.040504

Shannon, C. E. (1959). Coding theorems for a discrete source with a fidelity criterion. IRE National Convention Record, 7(4), 142-163.

Srednicki, M. (1993). Entropy and area. Physical Review Letters, 71, 666-669. https://doi.org/10.1103/PhysRevLett.71.666

Tishby, N., Pereira, F. C., & Bialek, W. (2000). The information bottleneck method. arXiv:physics/0004057.

Yao, Z., & Zhang, Z. (2019). Principal boundary on Riemannian manifolds. Journal of the American Statistical Association, 115(531), 1168-1181. Preprint: arXiv:1711.06705.

Suggested citation: Singh, R. (2026). The Entropy of Choice: Mathematical Appendices to the Boundary-Complexity Hypothesis. NSDM / MetaForgeX AI Working Paper 2026-02A, Revision 2.0. Johannesburg: MetaForgeX AI.