Channel Capacity of Domain-Specific Stochastic Resonance

Jeremy McEntire[1]

March 2026

Abstract

Papers VIII–IX measured domain selectivity as entropy change ratios. Paper XI established that geometric structure (effective dimensionality) bounds achievable selectivity. This paper completes the measurement chain by quantifying the information-theoretic content of domain-selective noise injection. We measure the KL divergence $D_{\mathrm{KL}}(P_\sigma \| P_0)$ between noisy and clean output distributions across a sigma sweep at the optimal injection layer (layer 10) from Paper IX.

All four domains exhibit inverted-U KL profiles, confirming stochastic resonance in information space. Total KL peaks at 14–15 bits ($\sigma^* = 0.2$–0.5), but the domain-*specific* component — the differential between self-domain and cross-domain KL — is small: +1.3 bits (medical), +1.9 bits (legal), −1.0 bits (code), −1.5 bits (science). Medical and legal show weak information selectivity (IS = 1.09–1.13); code and science are information-anti-selective (IS = 0.91–0.93). The domain-specific differential is consistent with the theoretical bound $C \leq \log_2(1 + k^2/d_{\mathrm{eff}}) = 2.24$ bits from Paper XI, resolving an apparent violation when total KL is naively compared to the bound.

1 Introduction

The stochastic resonance (SR) framework from Paper VI predicts that noise can benefit a suboptimal system when five sufficient conditions (C1–C5) are met. Papers VII and VIII tested this at the behavioral and activation levels respectively. But these measurements used entropy ratios and selectivity indices — surrogate metrics that do not directly quantify the *information content* of the noise signal.

This paper asks: how many bits of domain-specific information does shaped noise actually inject into the output distribution? The answer bounds the utility of any noise-based intervention: if the channel capacity is small, no sigma can produce large domain-selective effects. If it peaks at a nonzero sigma, SR is confirmed in information space.

We measure:

1. **KL divergence profile**: $D_{\mathrm{KL}}(P_\sigma \| P_0)$ as a function of σ

2. **Cross-domain KL matrix**: selectivity in information space

[1]Correspondence: `jmc@cageandmirror.com`

3. **Channel capacity**: $C = \max_\sigma \mathrm{MI}(\sigma) \geq \max_\sigma D_{\mathrm{KL}}$
4. **Inverted-U test**: does KL peak at interior σ, or is it monotonic?

2 Method

2.1 Setup

We inject domain-shaped noise at layer 10 (the optimal layer from Paper IX) using the `SingleLayerNoiseInjector` from Paper IX. The effective dimensionality at this layer is $d_{\mathrm{eff}} = 21.7$ (from Paper XI). The sigma sweep covers $\sigma \in \{0.01, 0.05, 0.1, 0.2, 0.5, 1.0\}$.

For each domain, 20 probes generate 32 tokens under both clean ($\sigma = 0$) and noisy ($\sigma > 0$) conditions. At each generation step, we record the full logit distribution.

2.2 KL Divergence Estimation

We approximate $D_{\mathrm{KL}}(P_\sigma \| P_0)$ using the union of top-100 tokens from both distributions at each generation step:

$$D_{\mathrm{KL}}(P_\sigma \| P_0) \approx \sum_{i \in S_P \cup S_Q} p_\sigma(i) \log_2 \frac{p_\sigma(i)}{p_0(i)} \tag{1}$$

This approximation is tight when the vocabulary is large and both distributions concentrate on few tokens (typical for greedy decoding).

2.3 Channel Capacity Bound

By the data processing inequality, $\mathrm{MI}(X;Y) \geq D_{\mathrm{KL}}(P_{Y|X=x} \| P_Y)$ for any specific x. Taking the max over σ:

$$C = \max_\sigma \mathrm{MI}(\sigma) \geq \max_\sigma D_{\mathrm{KL}}(P_\sigma \| P_0) \tag{2}$$

2.4 Theoretical Bound

Paper XI's concentration barrier bounds the INLP variance fraction by k/d_{eff}. We derive an approximate channel capacity bound by treating the domain-shaped perturbation as a signal in a Gaussian channel. If the k INLP directions per domain each contribute signal power proportional to their variance fraction, the aggregate signal-to-noise ratio scales as k^2/d_{eff} (summing k independent perturbation components each of magnitude k/d_{eff}). This motivates:

$$C_{\mathrm{domain}} \lesssim \log_2 \left(1 + \frac{k^2}{d_{\mathrm{eff}}}\right) \tag{3}$$

where $k = 9$ is the number of INLP directions per domain and d_{eff} is the effective dimensionality at the injection layer. With $d_{\text{eff}} = 21.7$, this gives $\log_2(1 + 81/21.7) = \log_2(4.73) = 2.24$ bits. This is an approximate upper bound, not a rigorous derivation; the primary finding is the empirical domain-specific differential below.

2.5 Information Selectivity

We define the information selectivity (IS) for domain d as:

$$\text{IS}_d = \frac{D_{\text{KL}}^{(d \to d)}}{D_{\text{KL}}^{(d \to \neg d)}} \tag{4}$$

where $D_{\text{KL}}^{(d \to d)}$ is the KL divergence measured on domain d probes when injecting domain d noise at optimal σ, and $D_{\text{KL}}^{(d \to \neg d)}$ is the mean KL on non-target domains. IS > 1 indicates domain-selective perturbation; IS < 1 indicates anti-selectivity.

3 Results

3.1 KL Divergence Profiles

Table 1 presents the mean KL divergence for each domain across the sigma sweep. All four domains exhibit inverted-U profiles.

Table 1: Mean $D_{\text{KL}}(P_\sigma \| P_0)$ in bits, by domain and σ. Bold: peak value per domain.

Domain	$\sigma = 0.01$	0.05	0.1	0.2	0.5	1.0
Medical	2.79	8.34	9.76	13.55	**14.55**	12.43
Legal	2.49	7.04	9.93	**15.12**	14.62	11.18
Code	2.65	7.84	9.48	**15.04**	14.60	11.36
Science	0.52	7.00	8.90	13.05	**14.03**	11.31

20 probes × 32 tokens per condition. σ is noise scale relative to hidden state norm.

Key observations:

1. KL rises steeply from $\sigma = 0.01$ to $\sigma = 0.2$, peaks at $\sigma = 0.2$–0.5, then declines at $\sigma = 1.0$.

2. Peak KL is ∼14–15 bits for all domains, reflecting substantial output distribution change.

3. Medical and science peak at $\sigma = 0.5$; legal and code peak at $\sigma = 0.2$.

4. Science starts lowest at $\sigma = 0.01$ (0.52 bits), suggesting science probes are less sensitive to weak perturbation.

3.2 Cross-Domain KL Matrix

Table 2 presents the KL divergence measured on each domain's probes when injecting noise shaped for each target domain, at the optimal sigma for each target.

Table 2: Cross-domain KL matrix at optimal σ^* for each target. Rows: noise target. Columns: measurement domain. Diagonal (self-KL) in bold.

Target ↓ / Meas. →	Medical	Legal	Code	Science
Medical ($\sigma^* = 0.5$)	**14.76**	14.53	14.05	11.89
Legal ($\sigma^* = 0.2$)	16.05	**16.63**	15.13	13.13
Code ($\sigma^* = 0.2$)	14.56	16.93	**13.85**	13.16
Science ($\sigma^* = 0.5$)	15.38	15.26	15.76	**14.00**

All values in bits. Computed at the sigma that maximizes target-domain KL.

The matrix is far from diagonal. Two domains exhibit information anti-selectivity:

- **Code-targeted noise** produces larger KL on legal (16.93) than on code (13.85). Code noise is *more informative* about legal outputs than code outputs.
- **Science-targeted noise** produces larger KL on code (15.76) and medical (15.38) than on science (14.00).

Medical and legal noise produce the largest KL on their own domains, but the margins are slim (+0.23 and +0.57 bits above the next-highest cross-domain value).

3.3 Channel Capacity and Information Selectivity

Table 3 summarizes per-domain channel capacity bounds and information selectivity.

The **total** channel capacity lower bound ($C_{\text{lb}} = 14$–15 bits) vastly exceeds the theoretical bound of 2.24 bits. This apparent violation is resolved by recognizing that the bound constrains domain-*specific* information, not total perturbation:

- The domain-specific differential Δ_d is +1.28 (medical), +1.86 (legal), −1.03 (code), −1.47 (science) — all within the 2.24-bit bound.
- The 14–15 bits of total KL are domain-*agnostic*: shaped noise changes the output distribution substantially but does so across all domains roughly equally.

Table 3: Per-domain channel capacity analysis. C_{lb}: channel capacity lower bound ($\max_\sigma D_{\text{KL}}$). IS: information selectivity. Δ_d: domain-specific differential (self-KL minus mean cross-KL). Bound: theoretical upper bound from d_{eff}.

Domain	C_{lb} (bits)	IS	Δ_d (bits)	Bound (bits)	Consistent?
Medical	14.55	1.095	+1.28	2.24	Yes
Legal	15.12	1.126	+1.86	2.24	Yes
Code	15.04	0.931	−1.03	2.24	Yes
Science	14.03	0.905	−1.47	2.24	Yes

$\Delta_d = D_{\text{KL}}^{(d\to d)} - \text{mean}(D_{\text{KL}}^{(d\to\neg d)})$. Bound: $\log_2(1 + k^2/d_{\text{eff}})$ with $k = 9$, $d_{\text{eff}} = 21.7$.

3.4 Inverted-U Test

All four domains exhibit inverted-U KL profiles, confirming SR in information space. The KL declines from its peak by 15–26% at $\sigma = 1.0$:

Table 4: Inverted-U evidence. σ^*: optimal sigma. Decline: percentage drop from peak to $\sigma = 1.0$.

Domain	σ^*	Decline at $\sigma = 1.0$	Shape
Medical	0.5	−15%	Inverted-U
Legal	0.2	−26%	Inverted-U
Code	0.2	−24%	Inverted-U
Science	0.5	−19%	Inverted-U

The inverted-U shape confirms Paper VI's prediction: there exists an optimal noise level beyond which additional perturbation destroys information rather than adding it. At $\sigma = 1.0$, the noise overwhelms the signal, reducing the KL divergence as the output distribution converges to a noise-dominated mode.

4 Discussion

4.1 Total vs. Domain-Specific Information

The central methodological insight is the distinction between total perturbation information and domain-specific information. Shaped noise at layer 10 injects ∼15 bits of information into the output distribution (enough to distinguish $2^{15} \approx 32{,}000$ output states). But nearly all of this information is shared across domains: noise shaped for code changes the medical output almost as much as the code output.

The domain-*specific* channel — the bits that distinguish the target domain from non-targets — carries only 1–2 bits for medical and legal, and is *negative* for code and science.

This is consistent with the concentration barrier: $k = 9$ INLP directions in $d_{\text{eff}} = 21.7$ effective dimensions can carry at most 2.24 bits of domain-specific information.

4.2 Domain Asymmetry in Information Space

The information selectivity pattern reproduces Paper IX's entropy selectivity pattern:

- **Medical and legal**: IS > 1 ($+1.3$ and $+1.9$ bits differential). These domains' INLP directions carry genuine, if small, domain-specific information content.
- **Code and science**: IS < 1 (-1.0 and -1.5 bits differential). Noise shaped for these domains is more informative about *other* domains' outputs.

This reproduces in information space the same asymmetry observed in entropy space (Paper IX) and validates the interpretation that code and science INLP directions target shared computational substrate rather than domain-specific pathways.

4.3 Information-Theoretic Interpretation of the Terminal Measurement Limit

Paper VIII's failure at terminal layers can now be restated in bits:

- Terminal injection produces high total KL (the output changes substantially) but near-zero Δ_d (the change is domain-agnostic).
- Paper IX showed intermediate layers (layer 10) improve selectivity modestly. The channel capacity measurement confirms: Δ_d at layer 10 is 1–2 bits for favorable domains, bounded by the concentration barrier.
- The concentration barrier caps Δ_d regardless of layer, but at terminal layers the large amplification factor (Paper X) ensures that the domain-agnostic component dominates even more completely.

The terminal measurement limit is an information-theoretic consequence of the concentration barrier: domain-specific bits are bounded by k/d_{eff} while total perturbation bits grow with σ, eventually making the domain-specific signal undetectable.

4.4 Connection to Paper VI

Paper VI's C1 condition (suboptimality) is the gate for SR. The inverted-U in KL space confirms C1 is satisfied: the clean model's output distribution is suboptimal in the sense that noise at moderate σ moves it to a state more distinguishable from baseline (higher KL). At

$\sigma = 1.0$, C1 is violated in a different sense — the noise exceeds the system's tolerance, and the output collapses toward a noise-dominated mode.

The optimal σ varies by domain (0.2 for legal/code, 0.5 for medical/science), consistent with Paper IX's finding that different domains have different sensitivity profiles.

5 Conclusion

We measured the channel capacity of domain-specific stochastic resonance at the optimal injection layer (layer 10) in Qwen-2.5 7B. The key findings:

1. All four domains exhibit inverted-U KL profiles, confirming SR in information space with optimal $\sigma = 0.2$–0.5.
2. Total channel capacity is large ($\sim$15 bits) but domain-agnostic.
3. Domain-specific information content is small: +1.3 to +1.9 bits for medical/legal, negative for code/science.
4. The domain-specific differential is consistent with the theoretical bound of 2.24 bits from Paper XI's concentration barrier.
5. Domain asymmetry (medical/legal selective, code/science anti-selective) is reproduced in information space, confirming Paper IX's entropy-based findings.

The information-theoretic picture completes the measurement chain: Paper VIII showed noise injection changes output distributions; Paper IX showed where selectivity peaks; Paper XI proved the geometric bound; Paper X showed the Jacobian provides no spectral shortcut. This paper quantifies the actual information content: domain-selective shaped noise carries at most $\sim$2 bits of domain-specific information per token, bounded by the concentration barrier. This is barely sufficient for a 4-way domain distinction and far too little for fine-grained domain-selective control.

References

[1] McEntire, J. (2026). Paper I: Leap+Verify. *arXiv:2602.19580*.

[2] McEntire, J. (2026). Paper II: Ensemble Collapse. *SSRN*.

[3] McEntire, J. (2026). Paper III: Constellation Composition.

[4] McEntire, J. (2026). Paper IV: Structural Transfer.

[5] McEntire, J. (2026). Paper V: Capability Manifold Surveillance.

[6] McEntire, J. (2026). Paper VI: Communicative Variance.

[7] McEntire, J. (2026). Paper VII: GenAI Is Socially Awkward.

[8] McEntire, J. (2026). Paper VIII: Shaped Noise Injection.

[9] McEntire, J. (2026). Paper IX: Layer-Resolved Response Tensor.

[10] McEntire, J. (2026). Paper X: Spectral Geometry.

[11] McEntire, J. (2026). Paper XI: The Concentration Barrier.