Opening B: adequate, not yet strong

Correct orientation; inadequate historical and conceptual positioning

For independent spins, single-spin logarithmic Sobolev inequalities tensorize. Interactions destroy the product structure, and the problem is to find assumptions under which a logarithmic Sobolev inequality remains uniform as the number of spins increases.

Bauerschmidt and Bodineau prove such an inequality for the $O(n)$ model under a condition on the spectral width of the interaction matrix. Their condition applies to the Sherrington–Kirkpatrick spin glass for $\beta < 1/4$. The argument introduces a Gaussian auxiliary field that renders the spins conditionally independent, and then uses a covariance bound to make the auxiliary-field measure uniformly log-concave.

What was already possible with earlier methods, and what is different here?

Opening C: strong (1 of 2)

Place the theorem in an existing mathematical problem

Logarithmic Sobolev inequalities provide quantitative control of relaxation to equilibrium for stochastic spin dynamics. For independent spins, such inequalities follow from one-spin estimates by tensorization. For interacting lattice spins, a substantial earlier theory established logarithmic Sobolev inequalities through forms of spatial mixing and control of conditional distributions; examples include work of Holley–Stroock and Stroock–Zegarlinski. These methods are effective when the dependence between sites can be controlled, but estimates based on absolute interaction strengths do not exploit cancellations between couplings of opposite signs.

This distinction becomes important for the Sherrington–Kirkpatrick spin glass. Its random interactions have both signs: their absolute row sums typically grow with the system size, whereas the spectral width of the interaction matrix remains bounded.

Context: Holley–Stroock (1987); Stroock–Zegarlinski (1992); Bauerschmidt–Bodineau (2018/2019).

Opening C: strong (2 of 2)

State the advance, its significance, and the proof mechanism

Bauerschmidt and Bodineau establish a uniform logarithmic Sobolev inequality for the $O(n)$ model whenever that spectral width is less than n. Applied to Ising spins with Sherrington–Kirkpatrick couplings, their theorem gives a uniform bound with probability tending to one for every fixed inverse temperature $\beta < 1/4$. This was presented by the authors as the first rapid-relaxation result for that model; the threshold is a sufficient condition, not a claimed optimal one.

The proof replaces a direct analysis of interactions by one Gaussian auxiliary-field representation. Conditional on this field, the spins are independent; a directional variance estimate makes the effective field potential uniformly convex, and an entropy decomposition transfers the resulting inequality back to the interacting system. We explain these steps, identify the one-spin results used as inputs, and show where the spectral condition enters.

Bauerschmidt–Bodineau, Theorem 1, Corollary 2, introductory discussion.

What does C accomplish that B does not?

1. **Historical baseline:** Earlier methods already established LSI in interacting systems under mixing assumptions.
2. **Precisely identified difficulty:** mixed-sign couplings create cancellations lost by absolute-value estimates.
3. **New mathematical advance:** a spectral-width condition with an application inaccessible to those earlier types of estimates.
4. **Proof architecture:** the reader learns not just what tool is used, but why it resolves the stated difficulty.

Context is a mathematical comparison, not decorative history.

Why compare proof passages?

mathematical setting, not a writing sample

Where we are in the BB proof

The entropy decomposition already controls the *conditional* spin energy. The remaining term is the energy of $G(\varphi) = \mu_\varphi(F^2)^{1/2}$ in the *auxiliary field*.

$$\nu_{\mathrm{r}}(|\nabla_{\varphi_x} G|^2) \quad \xrightarrow{\text{needed estimate}} \quad \nu(|\nabla_{\sigma_x} F|^2)$$

The writing assignment

Explain the same derivative/covariance estimate to an undergraduate reader. Compare Versions 2 and 3, two treatments of the same argument. Then consider the different organization of Version 4.

The goal is to assess expository choices, not to repeat the entire BB proof.

Version 2 — opening of the proof passage

CONSTRUCTED EXCERPT — opening paragraph and first calculation

Bauerschmidt and Bodineau prove a uniform logarithmic Sobolev inequality for the interacting $O(n)$ model. Their argument introduces an auxiliary Gaussian field φ such that the spins are independent conditional on φ. We now give an estimate used toward the end of their proof.

With $G(\varphi) = \mu_\varphi(F^2)^{1/2}$, differentiating the normalized expectation gives

$$\nabla_{\varphi_x} G = \frac{c}{2\mu_\varphi(F^2)^{1/2}} \operatorname{Cov}_{\mu_\varphi}(F^2, \sigma_x).$$

The covariance is a vector in $\mathbb{R}^n$. Assume F is not identically zero; otherwise the desired bound is trivial.

Constructed pedagogical example based on BB (2019), equations (9)–(15); see full comparison handout.

Version 3 — opening of the same passage

CONSTRUCTED EXCERPT — opening paragraph and first calculation

The Gaussian representation in Bauerschmidt–Bodineau writes the interacting $O(n)$ Gibbs measure as a *mixture* of product measures, not as a product measure itself. This matters after entropy decomposition: tensorization controls the conditional entropy, and the auxiliary-field LSI controls the entropy of $G(\varphi)^2$, but leaves the field derivative in the second term. To complete the theorem, we must bound this derivative using *spin* gradients.

The normalizing factor in the tilted one-spin law depends on φ_x, so differentiating gives a *covariance* rather than an uncentered moment:

$$\nabla_{\varphi_x} G = \frac{c}{2\mu_\varphi(F^2)^{1/2}} \operatorname{Cov}_{\mu_\varphi}(F^2, \sigma_x).$$

Constructed pedagogical example based on BB (2019), equations (9)–(15); see full comparison handout.

Version 2 — the covariance estimate

CONSTRUCTED EXCERPT — proof calculation; same mathematical setting

Fix the spins other than x and write $f(s) = F(\sigma_{\neq x}, s)$. If S, S' are independent with law μ_{φ_x}, then

$$\mathrm{Cov}(f(S)^2, S) = \tfrac{1}{2}\mathbb{E}[(f(S) - f(S'))(f(S) + f(S'))(S - S')].$$

By Cauchy–Schwarz, $|S - S'| \le 2$, and $\mathbb{E}(f(S) - f(S'))^2 = 2\,\mathrm{Var}(f(S))$, we obtain

$$|\,\mathrm{Cov}_{\mu_{\varphi_x}}(f^2, S)|^2 \le 8\,\mathrm{Var}_{\mu_{\varphi_x}}(f)\,\mu_{\varphi_x}(f^2).$$

Averaging over the remaining spins, using conditional independence and applying Cauchy–Schwarz a second time, gives the required full covariance estimate.

Constructed pedagogical example based on BB (2019), equations (9)–(15); see full comparison handout.

Version 3 — the covariance estimate

CONSTRUCTED EXCERPT — proof calculation; same mathematical setting

Hold the spins away from x fixed, write $f(s) = F(\sigma_{\neq x}, s)$, and take independent S, S' with law μ_{φ_x}. The duplication identity gives

$$\mathrm{Cov}(f(S)^2, S) = \tfrac{1}{2}\mathbb{E}[(f(S) - f(S'))(f(S) + f(S'))(S - S')].$$

This introduces $f(S) - f(S')$, which vanishes for constant f and is controlled by a Poincaré inequality. Since $|S - S'| \le 2$,

$$\mathbb{E}(f(S) - f(S'))^2 = 2\,\mathrm{Var}(f(S)), \quad \mathbb{E}(f(S) + f(S'))^2 \le 4\mathbb{E}f(S)^2,$$

so Cauchy–Schwarz gives the same local estimate as in Version 2. A *second* Cauchy–Schwarz inequality, now across the frozen spins, yields the full covariance estimate.

Constructed pedagogical example based on BB (2019), equations (9)–(15); see full comparison handout.

Compare Versions 2 and 3

analysis after the two specimens

Question: What explanatory work is the added prose doing?

- **Purpose:** it identifies the remaining *field-derivative term*, rather than saying only that an estimate appears "near the end."
- **Mathematical choice:** it explains why duplication is useful: the difference of F can be controlled by Poincaré.
- **Conditional structure:** it distinguishes a one-spin Cauchy–Schwarz inequality from the second application across the other spins.
- **Return to the theorem:** it explains why controlling that term closes the entropy argument, without adding a second proof.

More explanation is useful only where it clarifies the mathematical logic.

Version 4 — a different organization

CONSTRUCTED EXCERPT — selected claim and context

The final transfer estimate is logically independent of the spectral condition on the interaction matrix: that condition is used to obtain the auxiliary-field LSI. We therefore isolate the local calculation.

Claim. Let $\mu_\varphi = \bigotimes_x \mu_{\varphi_x}$, with $\mu_{\varphi_x}(ds) \propto e^{c\varphi_x \cdot s}\mu(ds)$ on S^{n-1}, and assume a uniform one-spin Poincaré inequality with constant γ^{-1}. For $G(\varphi) = \mu_\varphi(F^2)^{1/2}$,

$$\boxed{|\nabla_{\varphi_x} G|^2 \le \frac{2c^2}{\gamma}\mu_\varphi(|\nabla_{\sigma_x} F|^2).}$$

Application to BB. Combining this claim with the entropy decomposition gives

$$\mathrm{Ent}_\nu(F^2) \le \frac{2}{\gamma}\left(1 + \frac{2c^2}{\lambda}\right)\sum_x \nu(|\nabla_{\sigma_x} F|^2).$$

Constructed pedagogical example based on BB (2019), equations (9)–(15); see full comparison handout.

Is Version 4 necessarily better than Version 3?

the audience determines the answer

Question: What does isolating the claim accomplish?

- It separates a reusable conditional covariance argument from the spectral/convexity argument needed for the auxiliary-field LSI.
- It makes the exact assumptions and logical dependencies easier to locate and reuse.
- But it moves faster through the calculation; Version 3 may be better for classmates encountering the proof for the first time.

Conclusion

One might call Version 4 more "professional" in organization, but that does not make it the best version for every reader. Organization should serve the intended audience.

AI as a mathematical editor: a targeted request

a conditional workflow, not outsourced authorship

Example prompt: supply Version 2 and BB (9)–(15)

Propose **two brief additions** that would improve the reader's understanding: first, why the field-derivative term remains after entropy decomposition; second, why the proof uses Cauchy–Schwarz at two different stages. For each, supply at most two sentences and explain the mathematical benefit and the cost in length. Do not rewrite the proof or change any displayed estimate.

What makes the request productive?

The student first identifies the unresolved entropy term and the cited inputs; the model proposes *local alternatives with reasons*, leaving the final editorial decision to the student.

Suggested workflow; the course's actual AI-use rules govern what is authorized.

AI as a mathematical editor: judge the response

two valid ideas; only one is well placed here

Suggestion 1 — accept after checking

Add one sentence before differentiation: the only remaining uncontrolled entropy contribution involves $\nabla_{\varphi_x} G$, which must be estimated by spin gradients.

Suggestion 2 — usually reject here

Insert a full paragraph on the eigenvalues of the coupling matrix before deriving the local covariance bound.

Which serves this particular reader and this passage?

- **(1)** supplies exactly the mathematical purpose missing from Version 2.
- **(2)** supplies relevant mathematics, but duplicates earlier context and interrupts the local argument.

Illustrative suggestions, not an actual AI transcript.