Why the SK order parameter has no gaps

A variational argument from contact to strict curvature

Will Cook · Expository notes · 3 November 2026

AI assistance. Will Cook directed the work. AI agents based on large language models produced and revised this manuscript.

The Sherrington–Kirkpatrick (SK) model asks us to choose signs that maximise a random sum of pair interactions. Its limiting maximum is expressed by the Parisi formula through a nondecreasing function, the **order parameter**. We explain why a minimising order parameter is strictly increasing, that is, why it cannot be constant on any interval below one.

Suppose the order parameter were constant on an interval. Minimality fixes the values of an auxiliary curve at the two ends of that interval and bounds its slopes there, while an analytic estimate makes the derivative of the curve strictly convex in between. The curve is then left with too little area to connect the required values. Initial and terminal gaps need separate arguments. Finally we return to the spins, where a theorem about their energies gives nearly optimal sign choices differing on any prescribed positive fraction of positions up to one half, with arbitrarily small error in that fraction.

We follow the quantitative proof of OpenAI (2026a). Full support had already been proved by a different method in Hong-Bin Chen (2026), Theorem 1.3, and the estimate used here makes the final contradiction a direct comparison of areas. These notes assume algebra and elementary calculus, and they introduce probability and cumulative mass through calculations. Deeper analytic facts are stated where the proof needs them. By the end of Section 6 the reader should be able to reconstruct how minimality and those inputs exclude each kind of gap, and to translate the result into a statement about spin configurations. Section 7 collects optional calculations and the answers to the short questions in the text. The full analytic proofs are in the cited sources.

1. The finite optimisation problem

At each of n positions, or **sites**, we choose a sign $\sigma_i \in \{-1, 1\}$, called a **spin**, and a configuration is the vector $\sigma = (\sigma_1, \dots, \sigma_n)$. For each unordered pair we draw an independent standard Gaussian J_{ij}, a normal random variable with mean zero and variance one. Independence means that learning the other couplings leaves the probability distribution of this one unchanged. These coefficients, or **couplings**, are drawn once and then held fixed while we maximise the energy, or **Hamiltonian**,

$$H_n(\sigma) = \frac{1}{\sqrt{n}} \sum_{1 \le i < j \le n} J_{ij} \sigma_i \sigma_j.$$

A positive coupling rewards equal signs and a negative one rewards opposite signs. For example, positive couplings on pairs 12 and 23 favour $\sigma_1 = \sigma_2 = \sigma_3$, which conflicts with a negative coupling on pair 13. This inability to maximise every pair's contribution at once is called **frustration**. The factor $n^{-1/2}$ sets the scale of the model.

For two spins, $H_2(\sigma) = J_{12}\sigma_1\sigma_2/\sqrt{2}$, and choosing equal or opposite signs according to the sign of J_{12} gives $\max_\sigma H_2(\sigma) = |J_{12}|/\sqrt{2}$.

To average over the couplings, we first need their probabilities. A **probability density** assigns probability by area. For a standard Gaussian random variable Z, the density is $\varphi(z) = (2\pi)^{-1/2}e^{-z^2/2}$, which is nonnegative, has total integral one and, for $a \le b$, gives

$$\Pr(a \le Z \le b) = \int_a^b \varphi(z)\,dz.$$

In particular, a single point has probability zero, and the height of the density at a point is not itself a probability. To average a score $f(Z)$, we weight it by the same density, and its **expectation** is

$$\mathbb{E}f(Z) = \int_{\mathbb{R}} f(z)\varphi(z)\,dz.$$

For the standard Gaussian, $\mathbb{E}Z = 0$ and $\mathbb{E}Z^2 = 1$. The **variance** of a mean-zero random variable is its mean square, so $\sqrt{h}Z$ has variance h for $h > 0$. Symmetry gives the zero mean, whereas $\mathbb{E}|Z| = 2\int_0^\infty z\varphi(z)\,dz = \sqrt{2/\pi}$. The expected two-spin maximum is therefore $1/\sqrt{\pi}$. Signs chosen before the coupling is seen have expected energy zero, so the order matters: we choose signs for each realised coupling and then average the achieved maximum. For many spins we seek the limiting expected maximum **per spin**,

$$P_* := \lim_{n\to\infty} \frac{1}{n}\mathbb{E}\max_{\sigma} H_n(\sigma),$$

where the expectation averages over the couplings after the maximum is taken. **Zero temperature**, or **ground-state value**, refers here to this maximisation problem. The model is **pure** because every term contains exactly two spins, and has **zero external field** because there is no additional term $h\sum_i \sigma_i$.

Besides the maximum itself, we care about how the nearly optimal choices are arranged. To compare two configurations for the *same couplings*, we define their **overlap** $R(\sigma,\tau) = n^{-1}\sum_i \sigma_i\tau_i$, their average agreement with a disagreement counted as -1. If they disagree at k positions, then $R = 1-2k/n$. For example, $(+,+,+,+)$ and $(+,+,+,-)$ have overlap $1/2$, and two disagreements would give zero. We will ask whether two nearly optimal configurations can differ on roughly one quarter of their sites, and more generally at any prescribed fraction up to one half.

Our normalisation gives the covariance function $\xi(t) = t^2/2$ in the variational problem below. **Optional normalisation calculation.** For fixed σ,τ, put $v_i = \sigma_i\tau_i$. When we expand the product of energies, different couplings have mean product zero, by independence and their zero means, while a coupling multiplied by itself has mean one. Together with $v_i^2 = 1$, this gives

$$\mathbb{E}[H_n(\sigma)H_n(\tau)] = \frac{1}{n}\sum_{i<j} v_iv_j = \frac{(\sum_i v_i)^2 - n}{2n} = \frac{n}{2}R(\sigma,\tau)^2 - \frac{1}{2}.$$

Dividing by n gives the limiting covariance $\xi(t) = t^2/2$. For exact covariance $n\xi(R)$, add the same independent centred Gaussian of variance $1/2$ to every energy, which leaves the expected maximum unchanged.

Back to the spins. Reversing all the signs of one configuration preserves its energy and reverses the sign of its overlap with another. We next define deterministic candidate values indexed by nonnegative overlap scales t. The Parisi theorem of Section 3 identifies their minimum with the limiting spin optimum. That identification is an external theorem and involves no change of variables in the finite maximisation.

2. What a gap means

A candidate order parameter accumulates weight. Let us begin with a finite example, placing weight two at scale $1/4$ and weight one at scale $3/4$. The cumulative function is zero, then two, then three:

$$\gamma(t) = \begin{cases} 0, & 0 \le t < 1/4, \\ 2, & 1/4 \le t < 3/4, \\ 3, & 3/4 \le t < 1. \end{cases}$$

Its area is $2(1/2) + 3(1/4) = 7/4$. We can count the same area by following each weight: the first persists for three quarters of the interval and the second for one quarter, giving $2(3/4)+1(1/4) = 7/4$ again. Write δ_s for one unit of mass at s and $\mu = 2\delta_{1/4} + \delta_{3/4}$ for these weights. Integration against this measure is a weighted sum:

$$\int F(s)\,\mu(ds) = 2F(1/4) + F(3/4).$$

The weights total three, so this integral is not a probability average. We will reuse this accounting when each weight has a cost.

In general, an admissible **order parameter** is a function $\gamma : [0,1) \to [0,\infty)$ that is nondecreasing, right-continuous and has finite integral. Right-continuity means that at a jump the function already takes its new value, as in the example. The associated **Stieltjes measure** assigns to each interval the increment of γ across it:

$$\mu_\gamma([0,t]) = \gamma(t), \qquad \mu_\gamma((a,b]) = \gamma(b) - \gamma(a).$$

These cumulative values determine the measure uniquely. The first equation allows a mass $\gamma(0)$ at zero, and mass concentrated at one point is called an **atom**. These weights need not be probabilities of finite-system overlaps, and their total need not be one or even finite, although the mass of each $[0,T]$ with $T < 1$ is finite. Admissibility requires instead the weighted finiteness

$$\int_0^1 \gamma(t)\,dt = \int_{[0,1)} (1-s)\,\mu_\gamma(ds) < \infty.$$

This is the double count of area from the example: a unit at s contributes at every later time, for a total length $1-s$, and the general identity uses the theorem that nonnegative contributions may be summed in either order. A mass at one would change none of the cumulative values below one, which is why one is excluded from the domain.

The **support** consists of the points every neighbourhood of which carries positive mass, where a neighbourhood is an interval around the point, clipped to $[0, 1)$. Full support means that every nonempty relatively open interval carries mass. Equivalently, γ is strictly increasing: full support gives a positive increment on each $(a, b]$ with $a < b$, and a positive increment on a smaller interval inside any neighbourhood gives positive mass there.

For example, the admissible candidate

$$\gamma(t) = \begin{cases} t, & 0 \le t \le 1/3, \\ 1/3, & 1/3 \le t \le 2/3, \\ t - 1/3, & 2/3 \le t < 1 \end{cases}$$

has zero increment from $2/5$ to $3/5$, but increment $1/8$ from $1/8$ to $1/4$. Its measure has density one outside $(1/3, 2/3)$ and zero inside, in the sense that the mass of a set is the integral of this density over it. Both endpoints of the gap remain in the support, because every neighbourhood of either reaches a region carrying positive mass. The upper curve in Figure 1 stays positive across the gap while its *increments* vanish. This candidate has infinitely many support points, including points arbitrarily close to one, and still has a middle gap, so neither of those weaker properties gives full support.

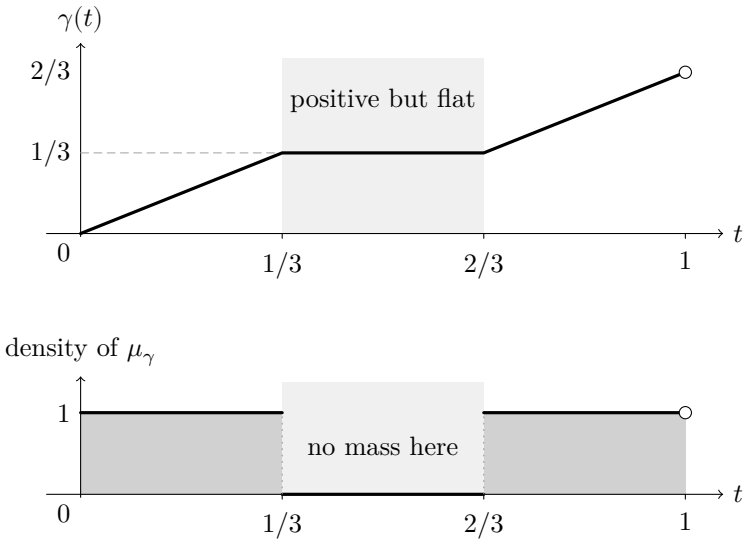

Figure 1: The same candidate in two representations. The upper curve is the piecewise-linear γ from Section 2; the lower plot shows the density of its Stieltjes measure. On $(1/3, 2/3)$ the cumulative value stays positive while its increments vanish. The endpoints $1/3$ and $2/3$ still belong to the support: every neighbourhood of either contains positive mass. The lower density belongs to this example; a general Stieltjes measure need not have one. Open circles mark the excluded endpoint 1. The candidate is not a computed minimiser. Axes and guides adapt OpenAI (2026a), Figure 1; the paired example is constructed here.

A constant candidate $\gamma(t) = c > 0$ sharpens the distinction. Its measure is $c\delta_0$, supported only at zero, so positivity of the function falls far short of the conclusion we seek. Full support does not require a density: summable positive masses at a dense countable set already give full support. The further-reading note returns to the stronger density conclusion.

Question 1: allowed changes. Return to the two weights $2\delta_{1/4} + \delta_{3/4}$. For $0 < \varepsilon < 1$, which changes leave an admissible cumulative function: adding ε mass at $1/2$, removing it at $1/4$, removing it at $1/2$, or multiplying both weights by $1-\varepsilon$? Check monotonicity as well as nonnegativity. Answer 1.

3. How a candidate receives its value

To rule out a candidate, we will look for an allowed change that lowers its value. Adding mass at one scale changes the cumulative coefficient at every later scale. We therefore need both a value for each candidate and a way to compute the cost of that change. The auxiliary optimisation below supplies the value, and its response to the starting state enters the first-variation formula of Section 5.1.

3.1 The zero coefficient

The variational problem assigns each γ a value by a one-dimensional auxiliary optimisation. Its final reward at state x is $|x|$, the best one-spin score $\max_{\sigma\in\{-1,1\}} x\sigma$. Let us first take $\gamma = 0$, so that only noise changes the state. The imported control representation uses noise with variance rate ξ'' (Auffinger–Chen, Theorem 2). Our normalisation $\xi(t) = t^2/2$ gives rate one, so the noise is **Brownian motion** B, a continuous random path starting at zero whose increments over disjoint time intervals are independent. Over an interval of length h, an increment has the law of the centred Gaussian $\sqrt{h}Z$, so from time t in state x the remaining noise $B_1 - B_t$ has the law of $\sqrt{1-t}Z$. The expected final reward is

$$\Phi_0(t,x) = \mathbb{E}|x + \sqrt{1-t}Z| = \int_{\mathbb{R}} |x + \sqrt{1-t}z|\varphi(z)\,dz.$$

At $(t,x) = (0,0)$ we recover the earlier calculation, $\Phi_0(0,0) = \sqrt{2/\pi}$, and at $t = 1$ no noise remains and the value is $|x|$. The function gives an expected payoff at a specified starting state, which should not be confused with a probability density.

3.2 Allowing a response to the noise

For a general γ, a response can bias the future state at a quadratic cost. The overlap scale t plays the role of time, and the path is an auxiliary object with no motion of spins behind it. Its noise is the Brownian motion introduced above.

A **control** $\alpha_s \in [-1,1]$ chooses a response at each time using only noise already observed since the starting time. The whole deterministic schedule γ is known in advance and the future noise is unknown, yet a rule fixed in advance can still respond differently on different observed paths. To see why the information constraint matters, let a sign $W = \pm 1$, each with probability $1/2$, be revealed after a single choice $v \in [-1,1]$. The expected score $|W+v| - v^2/2$ is $1 - v^2/2$, maximised at $v = 0$. The choice $v = W$ would score $3/2$, but it uses information that is unavailable when the choice is made. In the continuous control problem, decisions repeat under the same restriction, which the finite example only illustrates.

For a starting position x, the value is

$$\Phi_\gamma(t,x) = \sup_\alpha \mathbb{E}\left[\left|x + B_1 - B_t + \int_t^1 \gamma(s)\alpha_s\,ds\right| - \frac{1}{2}\int_t^1 \gamma(s)\alpha_s^2\,ds\right].$$

Inside the absolute value, noise and the accumulated bias change the state before its final reward is paid. The second integral is the running control cost. The supremum chooses the best permitted rule, and the formal version of the information condition is progressive measurability. Here the expectation averages over Brownian paths for a fixed γ, and the spin couplings play no part in it. At $\gamma = 0$ both control terms vanish and we recover Section 3.1.

The **Parisi functional** subtracts an explicit correction from this optimised value:

$$\mathcal{P}(\gamma) = \Phi_\gamma(0,0) - \frac{1}{2}\int_0^1 t\gamma(t)\,dt.$$

The inner supremum chooses a response for fixed γ, and the outer problem minimises $\mathcal{P}$ over γ. The ground-state Parisi formula and the attainment of its minimum give

$$P_* = \min_{\gamma \text{ admissible}} \mathcal{P}(\gamma).$$

Each candidate supplies an upper bound, and minimisation finds the sharpest one. This equality and attainment are external inputs from Auffinger–Chen (2017), Theorem 1 and the proof of Lemma 3, and Gaussian concentration gives the expected-value form used here. The no-gap theorem assumes a minimiser and proves a property of it, so the value $\mathcal{P}(0) = \sqrt{2/\pi}$ alone is silent on whether the zero candidate minimises.

3.3 The response, the state and the averaged square

We first regard the value as a function of the starting state, holding t and γ fixed, and write $\Phi = \Phi_\gamma$. Given a permitted control, use the same responses to each Brownian history at two starting states x and y. The costs are unchanged, and the triangle inequality bounds the change in the terminal reward by $|x - y|$, so taking the supremum over controls gives

$$|\Phi_\gamma(t,x) - \Phi_\gamma(t,y)| \le |x - y|.$$

The value is also **convex**: between two states its graph lies at or below the chord joining their values. Indeed, each fixed-control payoff is an average of translated absolute-value functions, hence convex. Its chord lies at or below the chord of the best endpoint payoffs, and taking the supremum preserves that upper bound.

OpenAI (2026a), Lemmas 2.1–2.2 supply spatial differentiability below time one. Let $u(t,x) = \partial_x\Phi_\gamma(t,x)$ be the slope as the starting state changes, so that a small shift δ changes the best expected payoff by approximately $u(t,x)\delta$, and write $a(t,x) = \partial_{xx}\Phi_\gamma(t,x)$ for the rate at which that slope changes. The displayed bound gives $|u| \le 1$, and convexity gives $a \ge 0$. Both are spatial derivatives, taken with time held fixed.

For the zero coefficient we can compute the slope directly. Put $F_{\mathrm{G}}(z) = \Pr(Z \le z) = \int_{-\infty}^{z} \varphi(y)\,dy$. Moving x slightly raises $|x + \sqrt{1-t}Z|$ when the final state is positive and lowers it when the final state is negative. Consequently

$$u_0(t,x) = 2F_{\mathrm{G}}\left(\frac{x}{\sqrt{1-t}}\right) - 1.$$

Differentiating under the integral is legitimate, because the difference quotients of the absolute value are bounded by one and the Gaussian hits the corner with probability zero. Thus the slope is the probability of a positive final state minus the probability of a negative one.

For general γ, the control theorem says that the slope response attains the value. To see why this is the right response, consider a short step of length h from (t,x). A response v adds displacement $h\gamma(t)v$, whose first-order contribution to the continuation value is $h\gamma(t)vu(t,x)$, at a running cost of $h\gamma(t)v^2/2$. The response-dependent gain per unit time is therefore

$$\gamma vu - \tfrac{1}{2}\gamma v^2 = \tfrac{1}{2}\gamma u^2 - \tfrac{1}{2}\gamma(v-u)^2.$$

Since $\gamma \ge 0$ and $|u| \le 1$, $v = u$ attains its maximum. The optimised contribution $\gamma u^2/2$ also explains why a squared slope will enter the sensitivity to γ. Completing the square is only a heuristic for differentiating an optimised value, and the control theorem and the first-variation formula of Section 5.1 make these connections rigorous. The control representation and feedback are justified for integrable coefficients in OpenAI (2026a), Lemmas 2.1–2.2; see also Auffinger–Chen, Corollary 2 and the extension after (13).

We first fix γ and obtain the deterministic value and slope by optimising the expected future payoff. We then use that slope as a rule along a random path. Computing an expected future payoff reveals nothing about the future noise, and at time t the rule needs only t and the state already reached. Following it generates the state

$$X_t = B_t + \int_0^t \gamma(s)u(s,X_s)\,ds, \qquad X_0 = 0.$$

This is the **associated diffusion**, whose existence and uniqueness for a fixed Brownian motion follow from the cited regularity theorem. The rule is called **feedback** because it uses the slope at the currently reached state. With γ and u supplied, a short forward step of length h from state x at time t, with $t + h < 1$, is approximated by

$$x + h\gamma(t)u(t,x) + \sqrt{h}Z.$$

The response is evaluated before the fresh independent noise Z is observed. This discrete approximation explains the order of operations, while the integral equation defines the continuous process. Its drift accumulates through time, so X stays continuous even where γ jumps.

Now fix γ, follow a path to X_t, evaluate the slope there, square it and average over paths. The result is a deterministic curve

$$q(t) := \mathbb{E}[u(t, X_t)^2].$$

Here Φ_γ and u are deterministic functions, while X_t and the response evaluated at it are random. Squaring before averaging keeps the size of both positive and negative responses.

A calculation at half time. For $\gamma = 0$ we have $X_t = B_t$, and $B_{1/2}$ has the same distribution as $Z/\sqrt{2}$. We can average over paths by grouping them according to the state reached, which at half time has density $e^{-x^2}/\sqrt{\pi}$. This density describes the **law** of the state, assigning probability to positions on the real line. The law has total mass one even though the order-parameter measure μ_0 is zero, because the two measures describe different objects.

We average $u_0(1/2, x)^2$ against this density. Substituting $x = z/\sqrt{2}$ gives

$$\begin{aligned} q_0(1/2) &= \int_{\mathbb{R}} (2F_{\mathrm{G}}(z) - 1)^2 \varphi(z)\, dz \\ &= \int_0^1 (2v - 1)^2\, dv = \frac{1}{3}, \qquad v = F_{\mathrm{G}}(z), \quad dv = \varphi(z)\, dz. \end{aligned}$$

The response at zero is zero, yet its mean square at the random state is $1/3$, so the identity $q(t) = t$ is far from automatic. Although every control is optimal at $\gamma = 0$, q_0 uses the value slope u_0. A discrepancy at one time is too little to lower $\mathcal{P}(0)$, and Section 7.1 uses the whole curve to find a feasible decrease.

Question 2: change the time and the future. Keep $\gamma = 0$ everywhere, but now set $t = 1/4$. What variance has the state accumulated, and what variance remains in its future reward? Write the Gaussian integral for $q_0(1/4)$ without evaluating it.

Now take $\gamma = 0$ before $1/2$ and $\gamma = 1$ afterward. Does that same integral still compute $q(1/4)$? Identify what remains unchanged and what must be recomputed. Answer 2.

The analytic facts needed in the proof are

$$q \in C^1([0, 1)), \qquad q(0) = 0, \qquad 0 \le q(t) \le 1, \qquad q'(t) = \mathbb{E}[a(t, X_t)^2].$$

Here C^1 means continuously differentiable, with derivatives at zero from the right. In particular, q' stays continuous across jumps of γ, and spatial curvature controls the growth of the squared response. These facts hold for every admissible candidate and follow from OpenAI (2026a), Lemmas 2.1–2.2.

The initial value has a direct symmetry explanation. Reflecting the starting state, the noise history and the response rule together preserves reward and cost, so $\Phi_\gamma(t, x)$ is even in x. Hence $u(t, 0) = 0$, and the zero-field starting state $X_0 = 0$ gives $q(0) = 0$. Section 7.2 explains the stochastic calculation for q'. In Section 5 the difference $q(t) - t$ will determine the first-order cost of changing cumulative weight, and since an added mass persists at every later time, that cost is a tail integral.

4. The full-support theorem

Source theorem (OpenAI (2026a), Theorem 1.1). In the pure, zero-field model with $\xi(t) = t^2/2$, let γ be an admissible integrable minimiser of $\mathcal{P}$. Its Stieltjes measure has full relative support on $[0, 1)$, and γ is strictly increasing. For every associated diffusion and every $0 \le t < 1$,

$$\mathbb{E}[u(t, X_t)^2] = t, \qquad \mathbb{E}[a(t, X_t)^2] = 1.$$

The hypothesis is minimality, since an arbitrary admissible candidate, such as the one in Figure 1, can have gaps. Zero field supplies $q(0) = 0$, and $\xi'' = 1$ enters the curvature estimate. Nonzero field and general mixed-spin covariance functions lie outside the statement.

5. Why a flat interval cannot minimise

We now prove the theorem, in which the identities $q = t$ and $q' = 1$ are conclusions still to be established. We begin with the first-order cost of an allowed change, derive conditions at support points, and then exclude every possible gap.

5.1 Allowed changes and their first-order cost

The **first variation** in a direction η is the derivative at zero of $\varepsilon \mapsto \mathcal{P}(\gamma + \varepsilon\eta)$. At a minimum its right derivative is nonnegative whenever small $\varepsilon \ge 0$ leave the candidate admissible. It need not vanish, since the opposite change may violate nonnegativity or monotonicity. Precisely, let $\eta \in L^1(0, 1)$, meaning $\int_0^1 |\eta(t)|\, dt < \infty$, and suppose $\gamma + \varepsilon\eta$ stays admissible for every $0 \le \varepsilon \le \varepsilon_0$, for some $\varepsilon_0 > 0$. The formula is

$$D\mathcal{P}(\gamma)[\eta] = \frac{1}{2} \int_0^1 \eta(t)(q(t) - t)\, dt.$$

Here q belongs to the unperturbed candidate γ. The term $-t/2$ comes from the outer correction in $\mathcal{P}$, and $q(t)/2$ is the sensitivity of the reoptimised control value. That sensitivity is the analytic content of the formula, which goes beyond differentiating with a fixed control. Chen–Handschy–Lerman (2018), Proposition 4 gives the derivative along the segment joining any two admissible coefficients. Taking the second endpoint to be $\gamma + \varepsilon_0\eta$ and rescaling the segment parameter gives the formula above. The detailed proof in OpenAI (2026a), Lemma 6.1 compares value functions and bounds the terminal remainder by the L^1 tail of η.

Adding a small atom at scale s raises the cumulative function at *every* later scale. The change is always allowed, because the total mass is free and no compensating removal is required. Its direction is $\eta(t) = \mathbf{1}_{[s,1)}(t)$, equal to zero before s and one from s onward. Thus

$$D\mathcal{P}(\gamma)[\mathbf{1}_{[s,1)}] = \frac{1}{2} \int_s^1 (q(t) - t)\, dt =: \frac{1}{2} G(s).$$

For the candidate in Figure 1, adding ε mass at $s = 1/2$ creates a jump of height ε there and raises the entire remaining curve by ε. The coefficient changes only on that tail, yet the optimised value and feedback must be recomputed, since earlier starting states also face a changed future control problem. The first-variation theorem accounts for this dependence, and its formula is a tail integral

because the tail is where the coefficient changes. Thus $G(s)/2$ is the marginal cost of placing mass at s, and at a minimiser $G(s) \geq 0$ for every $s < 1$.

For the two weights in Section 2, the same accounting gives

$$\int_0^1 \gamma(t)(q(t) - t)\, dt = 2G(1/4) + G(3/4).$$

Each weight is charged for its entire remaining tail. The q and G in this identity belong to that candidate, and so far no optimality has been assumed. Reducing both weights by a fraction ε multiplies the cumulative function by $1 - \varepsilon$.

Suppose temporarily that this two-weight candidate minimised $\mathcal{P}$. Addition would give two nonnegative costs, $G(1/4) \geq 0$ and $G(3/4) \geq 0$, while the allowed reduction would give

$$0 \leq -\frac{1}{2}(2G(1/4) + G(3/4)).$$

Both costs must therefore be zero, since their positive-weighted sum is nonnegative and the reduction forces it to be nonpositive. This is a necessary condition under the temporary minimiser assumption, and the two-weight candidate need not attain the minimum.

Now fix a general minimiser γ and its own q, G and μ_γ. The same argument works with the finite sum replaced by an integral. The direction $\eta = -\gamma$ is integrable, and $(1 - \varepsilon)\gamma$ remains admissible for $0 \leq \varepsilon \leq 1$, scaling all existing mass by the same factor. For the continuous candidate in Figure 1 it multiplies the two nonzero slopes and the plateau height by $1 - \varepsilon$, so it works for diffuse mass as well as atoms. To see the order of integration, note that each weight at s contributes for exactly the times $t \geq s$:

$$\begin{aligned}\int_0^1 \gamma(t)(q(t) - t)\, dt &= \int_0^1 \int_{[0,1)} \mathbf{1}_{\{s \leq t\}}(q(t) - t)\, \mu_\gamma(ds)\, dt \\ &= \int_{[0,1)} G(s)\, \mu_\gamma(ds).\end{aligned}$$

Thus the allowed reduction gives

$$0 \leq D\mathcal{P}(\gamma)[-\gamma] = -\frac{1}{2}\int_{[0,1)} G(s)\, \mu_\gamma(ds).$$

The exchange of integration order is valid even if the total mass is infinite, because the absolute integral is bounded by the area from Section 2:

$$\int_{[0,1)} \int_s^1 |q(t) - t|\, dt\, \mu_\gamma(ds) = \int_0^1 \gamma(t)|q(t) - t|\, dt \leq \int_0^1 \gamma(t)\, dt < \infty.$$

Since $G \geq 0$, its integral must be zero. If G were positive at a support point, continuity would bound it below by a positive number on a neighbourhood carrying positive mass, contradicting the zero integral. Hence $G = 0$ on the support. We call this **contact**: the nonnegative function touches

zero there. The scaling argument is the variational argument of Chen, Proposition 3.2, equations (3.11)–(3.12), applied with the integrable-direction formula above.

Scaling removes only mass that exists. In Figure 1 the supported point $1/4$ carries no atom, since intervals around it carry mass while the point itself does not. Subtracting $\varepsilon\delta_{1/4}$ would create a downward jump, and for $0 < \varepsilon < 1/4$ the cumulative values would remain nonnegative while the function stopped being nondecreasing. Support alone therefore permits no atom removal. What we have proved is contact on the support, and the argument allows G to vanish at points that carry no mass.

5.2 Contact gives endpoint conditions

We need information about q at the ends of a possible gap. An interior support point s is a minimum of G. Since $G'(s) = s - q(s)$ and $G''(s) = 1 - q'(s)$, ordinary calculus gives

$$q(s) = s, \qquad q'(s) \le 1.$$

If zero belongs to the support, contact and symmetry give $G(0) = q(0) = 0$. Were $q'(0) > 1$, continuity would give $q'(t) > 1$ for small positive t, and integration would then give $q(t) > t$, so $G'(t) = t - q(t) < 0$. Starting from $G(0) = 0$, this would force $G(t) < 0$, a contradiction. Hence $q'(0) \le 1$, and this one-sided argument covers an atom at zero.

5.3 Strict curvature creates an impossible integral

The support is relatively closed in $[0, 1)$. Starting at a missing point, enlarge its mass-free interval until it reaches support or a boundary of the domain. Such a maximal interval has two support endpoints, starts at zero outside the support, or reaches one. We treat these three cases in order, and the last includes empty support.

First suppose (r, s) has both endpoints in the support, with $0 \le r < s < 1$. Its constant height c is necessarily positive: for t inside the gap, a small relative neighbourhood of r lies in $[0, t]$ and has positive mass, so that $c = \mu_\gamma([0, t]) > 0$. If $r = 0$, this mass is an atom at zero. Contact imposes

$$q(r) = r, \quad q(s) = s, \quad q'(r) \le 1, \quad q'(s) \le 1.$$

The first two conditions require area $s - r$ under q' across the gap, and a calculus model shows how strict convexity can prevent this. On the earlier interval $[1/3, 2/3]$, put $x = 3t - 1$ and consider the slope $1 - x(1 - x)$. Its endpoints equal one, but

$$\int_{1/3}^{2/3} [1 - (3t - 1)(2 - 3t)]\, dt = \frac{1}{3} \int_0^1 [1 - x(1 - x)]\, dx = \frac{5}{18}.$$

The missing area is $1/3 - 5/18 = 1/18$. The model is a calculus example chosen for its shape, and the next input forces the actual q' below its endpoint chord in the same way.

Strict-curvature input (OpenAI (2026a), Proposition 5.1). For any admissible coefficient equal to $c > 0$ on an open interval, its moment curve satisfies, strictly inside that interval,

$$q'''(t) \geq \frac{c^2}{1000}\mathbb{E}[a(t, X_t)^4] > 0 \qquad (r < t < s).$$

Section 7.3 explains how this sign is obtained after expectation. The gap argument uses only the strict sign, whatever the constant in place of $1/1000$. The final inequality uses $a(t, x) > 0$, proved by Gaussian smoothing on the constant interval (OpenAI (2026a), Lemma 2.3), whereas convexity of Φ alone would give only $a \geq 0$.

Because $q''' = (q')''$, the function q' is strictly convex: between two points its graph lies strictly below their joining chord. The moment identity makes q' continuous at the endpoints, even if γ jumps there, and higher derivatives are needed only inside the gap. Hence

$$q'(t) < \frac{s-t}{s-r}q'(r) + \frac{t-r}{s-r}q'(s) \leq 1 \qquad (r < t < s).$$

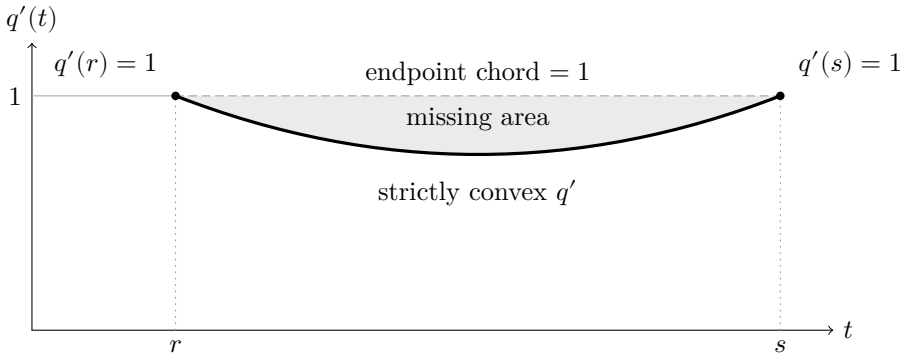

Figure 2: The case where strictness is indispensable: both endpoint slopes equal one, so the chord alone loses no area. A strictly convex slope lies below it. Contact would require zero shaded area, but $\int_r^s (1 - q'(t))\,dt > 0$. The drawn calculus model is $1 - x(1-x)$ with $x = (t-r)/(s-r)$; its mean is $5/6$, while the constant slope one has mean one. Neither is asserted to be an SK moment curve. The proof also allows unequal endpoint slopes, each at most one. Adapted from Figure 1 of OpenAI (2026a), OpenAI/math (Apache 2.0).

Integration gives $q(s) - q(r) < s - r$, contradicting contact. For the candidate of Section 2, a minimising plateau on $(1/3, 2/3)$ would require the same increment to be both equal to and strictly less than $1/3$.

5.4 The remaining initial and terminal gaps

Initial gap with zero outside the support. Suppose $[0, s)$ is missing and $s < 1$ is the first support point. Then $\gamma = 0$ there, so the positive-coefficient bound gives no strictness, and the zero-coefficient heat identity takes its place:

$$q''(t) = \mathbb{E}[\Phi_{xxx}(t, X_t)^2] > 0 \qquad (0 < t < s).$$

The inequality is strict for the following reason. On this interval $X_t = B_t$, whose density is positive everywhere at positive times. If the expected square were zero, continuity would force

$\Phi_{xxx}(t,x) = 0$ for every x, since a nonzero value would give a small interval of positive probability on which the square is positive. Then $\Phi(t,\cdot)$ would be a quadratic polynomial, and its bounded slope would make it affine, a linear function plus a constant. Choosing zero control, however, gives $\Phi_\gamma(t,x) \geq \mathbb{E}|x + B_1 - B_t| \geq |x|$, because the average absolute value is at least the absolute value of the average. An affine function cannot dominate $|x|$ in this way: the positive tail would require slope at least one and the negative tail slope at most minus one. The heat identity and the needed regularity are in OpenAI (2026a), Lemmas 2.2–2.3.

Thus q' is strictly increasing, and right contact gives $q'(t) < q'(s) \leq 1$ for $0 < t < s$. Symmetry supplies $q(0) = 0$ and contact gives $q(s) = s$, so integrating again leaves too little area. Strictness was needed only at positive times, which matters because symmetry can make the third spatial derivative vanish at $(0,0)$.

Although the forward path is Brownian here, the backward value receives the data $\Phi_\gamma(s,\cdot)$ from the later coefficient, so the formula for the globally zero candidate cannot be reused on this initial interval. The case in which zero itself is supported was already covered in Section 5.3, using its atom and the one-sided slope bound.

Empty support. Empty support would mean $\gamma = 0$ everywhere. The preceding heat argument then gives $q'' > 0$ throughout $(0,1)$, so q itself is strictly convex. Recall that $q(0) = 0$ and $q \leq 1$. For $0 < t < s < 1$, the chord from zero to s gives

$$q(t) \leq \frac{t}{s} q(s) \leq \frac{t}{s}.$$

Letting s increase to one gives $q(t) \leq t$, where only the numerical bound t/s is taken to a limit. For strictness we choose $t < v < 1$ and compare with the chord ending at v:

$$q(t) < \frac{t}{v} q(v) \leq t.$$

Hence $G(0) = \int_0^1 (q(t) - t)\, dt < 0$, contrary to the allowed addition of mass at zero. The argument uses no value or derivative at one. Section 7.1 also computes the exact first variation $D\mathcal{P}(0)[1] = 1/4 - 1/\pi$.

Terminal gap. Now suppose a maximal missing interval is $(r,1)$. Its left endpoint $r < 1$ is supported, and its constant height is positive, by the same mass argument as in Section 5.3. We have $q(r) = r$ and $q'(r) \leq 1$, and q' is strictly convex on $(r,1)$. With no contact condition at one, we use instead the bound $q \leq 1$, which holds because $q = \mathbb{E}[u^2]$ and $|u| \leq 1$. We show in two steps that q stays below the diagonal throughout the gap, which makes the tail cost $G(r)$ negative.

Reaching the diagonal forces a slope above one. Put $h(t) = q'(t) - 1$, so $h(r) \leq 0$. If $q(t) \geq t$ at some $r < t < 1$, strict integration below the chord on the finite interval $[r,t]$ would give

$$0 \leq q(t) - t = \int_r^t h(v)\, dv < \frac{t-r}{2}\big(h(r) + h(t)\big) \leq \frac{t-r}{2} h(t).$$

Hence $h(t) > 0$.

The excess slope persists and violates the bound. To control a later slope, we apply the same chord rule on the larger interval $[r,s]$ with $t < s < 1$. At its interior point t this gives

$$h(t) \le \frac{s-t}{s-r}h(r) + \frac{t-r}{s-r}h(s).$$

Solving for the right endpoint gives

$$h(s) \ge h(t) + \frac{s-t}{t-r}(h(t) - h(r)) \ge h(t) > 0.$$

The bound holds at every time after t, so integration gives $q(s) \ge t + (s-t)(1+h(t))$. Choose any time

$$t + \frac{1-t}{1+h(t)} < s < 1.$$

Such a choice exists because $h(t) > 0$, and the lower bound now gives $q(s) > 1$, contradicting $q(s) \le 1$ strictly before time one. Thus $q(t) < t$ throughout $(r, 1)$, giving $G(r) = \int_r^1 (q(t) - t)\,dt < 0$, contrary to the addition test. The argument uses only finite-interval convexity and the bound on q, and it assumes no value for $q'(1)$.

By a different terminal argument, the source proves the stronger estimate $1 - q(t) \ge C\sqrt{1-t}$ near one for any finite constant terminal coefficient, with no need for the left contact data used here. The proof is OpenAI (2026a), Lemma 6.3, reproduced in the *Long Reasoning Record*, §5, terminal estimate. The absence of terminal plateaus had already been noted in Auffinger–Chen–Zeng, Remark 7.

These cases exhaust the complement of the support. Contact now gives $q(t) = t$ for every interior time, and symmetry adds the time zero. Differentiating gives $\mathbb{E}[a(t, X_t)^2] = 1$, which completes the proof of the theorem.

5.5 Two checks on the gap argument

Question 3: tilt the chord. Suppose a differentiable function f on $[2, 5]$ has $f(2) = 2$, $f(5) = 5$, $f'(2) = 1/2$ and $f'(5) = 3/2$. If f' is strictly convex, can these data hold? What changes if strict convexity is weakened to convexity? The endpoint chord is now tilted, so a pointwise comparison with the constant one is unavailable. Finally, keep the endpoint values and left slope fixed but change $f'(5)$ to 2. Are these data compatible with strict convexity of f'?

Question 4: the missing right endpoint. Put the proof aside and explain why a terminal plateau cannot minimise. Identify what replaces the missing right contact point and why the contradiction occurs before time one. Then check whether the exact bound $q \le 1$ could be replaced by mere boundedness: test the calculus function $q(t) = t - 3t^2/4 + t^3$ on $[0, 1)$ against the remaining contact and curvature conditions at $r = 0$. The test concerns the calculus argument only, and the polynomial serves as a calculus example.

Answers 3 and 4.

6. Return to the spins

At a fixed finite size, almost surely the only exact maximisers are a configuration and its global reversal. Indeed, two configurations that are neither equal nor global reversals differ in some pair product. Their energy difference is then a Gaussian with positive variance, so a tie has probability zero, and there are only finitely many pairs. Intermediate separations appear once we allow a small loss per spin and let the system grow.

A separate theorem of Auffinger–Chen–Zeng, Proposition 1, equation (5) applies at each point of the minimiser's support, and combining it with full support gives the following. For every fixed overlap $r \in [0, 1)$ and tolerances $\varepsilon, \eta > 0$, with probability over the couplings tending to one exponentially fast as $n \to \infty$, there exist configurations σ, τ with

$$|R(\sigma, \tau) - r| < \varepsilon, \qquad H_n(\sigma)/n,\ H_n(\tau)/n \geq P_* - \eta.$$

If k sites differ, $k/n = (1 - R)/2$. The overlap tolerance therefore gives

$$\left|\frac{k}{n} - \frac{1 - r}{2}\right| < \frac{\varepsilon}{2}.$$

For example, fix an energy loss $\eta > 0$ per spin and take $r = 1/2$ and $\varepsilon = 0.02$. With probability tending to one exponentially fast, the same instance has two configurations with energies per spin at least $P_* - \eta$ and a differing-site fraction between 0.24 and 0.26. No numerical threshold for n is asserted. At target zero the differing-site fraction can be arbitrarily close to one half.

The pair may depend on the chosen target and on the couplings. The theorem asserts existence for each fixed target and says nothing about typically sampled pairs or about how to find them. The companion *Long Reasoning Record*, §9.3, "Near-ground configurations at a prescribed overlap", gives the probability bound, covariance adjustment and exact finite-size comparison. The threshold $P_* - \eta$ refers to the limiting value, whereas an exact maximiser attains its particular finite instance's maximum.

Question 5: use the spin conclusion. Fix an energy loss $\eta > 0$. How can the theorem give two configurations differing on between 69% and 71% of sites, although it is stated for nonnegative overlap targets? Which earlier symmetry is needed? Does the stated theorem also permit the choice $\eta = 1/n$? Answer 5.

7. Optional calculations

With the no-gap argument and the return to spins complete, this section collects optional calculations. Section 7.1 confirms the zero-candidate descent by Gaussian calculus, using the stated moment identity, and Section 7.2 explains the value equation and stochastic chain rule behind that identity. Section 7.3 describes the mechanism of the strict-curvature estimate, whose full proof is in the sources, and Section 7.4 answers the earlier questions.

7.1 A descent direction for the zero candidate

We return to the Gaussian value and slope of Section 3. With $\gamma = 0$ the state is B_t, and differentiating the slope once more gives

$$X_t = B_t, \qquad u_0(t,x) = 2F_{\mathrm{G}}\left(\frac{x}{\sqrt{1-t}}\right) - 1, \qquad a_0(t,x) = \sqrt{\frac{2}{\pi(1-t)}}e^{-x^2/(2(1-t))}.$$

A Gaussian integral gives $\mathbb{E}e^{-B_t^2/(1-t)} = \sqrt{(1-t)/(1+t)}$. Substituting into $q_0' = \mathbb{E}[a_0(t,B_t)^2]$ and integrating from zero yields

$$q_0'(t) = \frac{2}{\pi\sqrt{1-t^2}}, \qquad q_0(t) = \frac{2}{\pi}\arcsin t.$$

This recovers the half-time value $1/3$, and it shows that $q_0(t) \to 1$ while $q_0'(t) \to \infty$, so convergence of a function at one gives no bound on its derivative there. The constant direction $\eta = 1$ is admissible and adds an atom at zero, at the cost

$$D\mathcal{P}(0)[1] = \frac{1}{2}\int_0^1 \left(\frac{2}{\pi}\arcsin t - t\right) dt = \frac{1}{4} - \frac{1}{\pi} < 0.$$

A small positive constant coefficient therefore improves the zero candidate. The calculation uses the heat and variation formulas and no minimiser hypothesis (see also Chen–Handschy–Lerman, proof of Theorem 6).

7.2 The value equation and the moment identity

At $t = 1$ there is no noise or response left to accumulate, so the value is $|x|$. At earlier times it includes the best remaining response. Its partial differential equation (PDE) runs backward from that final value:

$$\partial_t \Phi_\gamma = -\frac{1}{2}\left(\partial_{xx}\Phi_\gamma + \gamma(t)(\partial_x \Phi_\gamma)^2\right), \qquad \Phi_\gamma(1,x) = |x|.$$

With the spatial derivatives u and a defined in Section 3.3, the equation expresses a recursive choice: choose a response for a short first step, then use the best continuation from the state reached. At a continuity point of γ, this gives the formal first-order relation

$$\Phi_\gamma(t,x) = \sup_{|v|\le 1}\left\{\mathbb{E}\Phi_\gamma(t+h, x+\gamma(t)vh + \sqrt{h}Z) - \tfrac{1}{2}\gamma(t)v^2 h\right\} + o(h).$$

Subtracting the present value, dividing by h and expanding by Taylor's formula, with $\mathbb{E}Z = 0$ and $\mathbb{E}Z^2 = 1$, we obtain

$$0 = \Phi_t + \tfrac{1}{2}\Phi_{xx} + \sup_{|v|\le 1}\gamma(t)(v\Phi_x - v^2/2).$$

The square completion from Section 3.3 gives a supremum of $\gamma u^2/2$, attained at $v = u$, and hence the stated PDE. The calculation explains the equation, while justifying the recursion and the Taylor expansion requires the control representation and the regularity results. See Auffinger–Chen, Theorem 2 and Corollary 2 for the zero-temperature control formula and Jagannath–Tobasco, Lemma 18 for the control method with finite-temperature terminal data. The extension after Auffinger–Chen's Corollary 2 covers integrable coefficients. The regularity used here is in OpenAI (2026a), Lemmas 2.1–2.2.

Itô's rule is the chain rule for Brownian paths. Along $dX_t = b_t\,dt + dB_t$, the coefficient of dt (the **drift**) in a smooth $f(t, X_t)$ is $f_t + b_t f_x + f_{xx}/2$, and its noise term is $f_x\,dB_t$. If we differentiate the PDE in x and take $b_t = \gamma u$, the drift of the slope cancels:

$$du(t, X_t) = \underbrace{(u_t + \gamma u u_x + u_{xx}/2)}_{0}\,dt + a(t, X_t)\,dB_t.$$

Thus $M_t = u(t, X_t)$ is a bounded **martingale**: its conditional expected later value, given observations up to time t, is its present value. Zero field gives symmetry and $M_0 = u(0, 0) = 0$. Applying Itô's rule once more to M_t^2 gives

$$d(M_t^2) = 2M_t a(t, X_t)\,dB_t + a(t, X_t)^2\,dt.$$

The noise term has mean zero because its coefficient is nonanticipating and bounded on each strip ending before one. Taking expectations and integrating yields $q(t) = \int_0^t \mathbb{E}[a(s, X_s)^2]\,ds$, which gives the asserted derivative. Since $\mathbb{E}M_t = 0$, $q(t)$ is also the variance of the slope response.

For each fixed $T < 1$, the required spatial derivatives are bounded on $[0, T] \times \mathbb{R}$, and step-coefficient and smooth-terminal approximations justify the identities. The bounds depend on the coefficient and on the strip, which ends before time one, and the PDE holds for almost every time.

7.3 Where the strict-curvature input comes from

The proof in Section 5 used strict curvature on a positive constant interval. We now explain how a correction with expectation zero makes the estimate usable once the required density bounds and algebraic identity are supplied, both of which the sources construct and prove. Take any admissible coefficient, minimising or not, with $\gamma = c > 0$ on an open interval, and set $b = \Phi_{xxx}$ and $d = \Phi_{xxxx}$ there. Differentiating the PDE and applying Itô's rule gives

$$q'' = \mathbb{E}[b^2 - 2ca^3], \qquad q''' = \mathbb{E}[d^2 - 12cab^2 + 6c^2a^4].$$

The negative middle term rules out a pointwise positivity argument, so the source averages instead with the spatial probability density $f_t(x)$ of X_t at positive time. As in the half-time example, grouping paths by their reached position gives, for any integrable function F,

$$\mathbb{E}[F(X_t)] = \int_{\mathbb{R}} F(x) f_t(x)\,dx.$$

These probability weights live on states x and are distinct from the Stieltjes weights $\mu_\gamma(ds)$ on overlap scales. The proof needs bounds relating this density to the spatial derivatives of the value,

and OpenAI (2026a), §3 gives the factorisation and comparison arguments that establish them. An earlier comparison between the density and the value's derivatives is in Zhou (2026), Appendix E.3. Section 4.3 of the long record shows how Zhou's inequalities give this density comparison.

The operation can be written directly for the positive density f_t, where f_t' means differentiation in the state variable x with t fixed. For a smooth function F with integrable terms and vanishing boundary product Ff_t,

$$\mathbb{E}\left[F'(X_t) + F(X_t)\frac{f_t'(X_t)}{f_t(X_t)}\right] = \int_{\mathbb{R}} (Ff_t)'(x)\,dx = 0.$$

For a small example of this operation, let Z be standard normal and $F(x) = 2x$. Its density has logarithmic derivative $-x$, so $\mathbb{E}[2 - 2Z^2] = 0$. Although $2x^2 - 1$ is negative near zero, adding that zero-mean correction gives

$$(2x^2 - 1) + (2 - 2x^2) = 1.$$

Hence $\mathbb{E}[2Z^2 - 1] = 1 > 0$. The boundary term vanishes because a polynomial times the Gaussian density tends to zero. This illustrates how a correction changes the visible sign without changing the expectation, and it supplies none of the SK-specific sign estimates.

The source first works with positive nondecreasing finite-step coefficients and the terminal approximation $M^{-1}\log(2\cosh(Mx))$, with M at least as large as every step height, and its sign comparisons use this bound as well as smoothness. At this stage write $R_c = d^2 - 12cab^2 + 6c^2a^4$. The source constructs a correction C from the spatial derivatives and the density, with

$$\mathbb{E}[C(t, X_t)] = 0, \qquad R_c(t, x) + C(t, x) \geq \frac{c^2}{1000}a(t, x)^4.$$

A rational polynomial identity gives the second inequality by writing the remainder as a sum of nonnegative terms. The density bounds and vanishing boundary products justify the first. Consequently

$$q'''(t) = \mathbb{E}[R_c(t, X_t)] = \mathbb{E}[(R_c + C)(t, X_t)] \geq \frac{c^2}{1000}\mathbb{E}[a(t, X_t)^4].$$

This is how the correction produces a useful lower bound for the expectation even though R_c can be negative at particular states.

To reach an arbitrary admissible coefficient, fix a time inside the plateau and approximate in L^1 by step coefficients, keeping the same value c near that time. Uniform convergence of spatial derivatives through order four on a fixed strip below one, together with convergence of the coupled paths, passes the quantitative inequality to the limit. Strictness is then established for the limit itself: the constant-coefficient heat formula expresses a as a nonnegative mean plus c times a strictly positive variance of a nonconstant slope under a positively weighted Gaussian law. Hence $a(t, x) > 0$ and its fourth moment is positive. Positivity of the approximants alone would not suffice, since it could be lost in the limit. The auxiliary logarithmic derivatives need not converge, because integration has already removed them.

The analytic details are in OpenAI (2026a), §§3–5 and in the companion *Long Reasoning Record*, §4.4, "The polynomial correction with zero expectation", and §4.5, "Coercivity at a fixed time and on plateaus".

7.4 Answers to the questions

Question 1. Adding mass at $1/2$ gives cumulative heights $0, 2, 2+\varepsilon, 3+\varepsilon$, in their successive intervals; removing it at $1/4$ gives heights $0, 2-\varepsilon, 3-\varepsilon$. Both functions remain nonnegative and nondecreasing. Removing mass at $1/2$ instead lowers the height from 2 to $2-\varepsilon$ there, creating a downward jump, so it is forbidden even though all values remain nonnegative. Scaling both weights preserves all admissibility conditions. These tests concern which candidates are allowed, and a smaller ordinary area alone gives no comparison of Parisi values.

Question 2. For $\gamma = 0$, the state is $X_{1/4} = B_{1/4}$, a centred Gaussian with variance $1/4$, so it has the law of $Z/2$. The remaining independent Gaussian noise $B_1 - B_{1/4}$ has variance $3/4$ and standard deviation $\sqrt{3}/2$. Thus

$$q_0(1/4) = \int_{\mathbb{R}} \left(2F_{\mathrm{G}}(z/\sqrt{3}) - 1\right)^2 \varphi(z)\, dz.$$

The two variances agree at half time and differ here. For nonzero z the response is now smaller in magnitude, so $q_0(1/4) < 1/3$, as the arcsine formula of Section 7.1 confirms.

Changing the coefficient after $1/2$ leaves $X_{1/4} = B_{1/4}$, because its drift has been zero up to that time. The remaining Brownian noise still has variance $3/4$, but the value now includes future control terms, and its slope must be recomputed before the square is averaged against the unchanged state law. The displayed integral using u_0 therefore no longer applies. Knowing the future coefficient schedule reveals nothing about the future random noise.

Question 3. A strictly convex f' lies below its endpoint chord. That chord has mean height $(1/2 + 3/2)/2 = 1$, hence integral three. Therefore $f(5) - f(2) = \int_2^5 f'(t)\, dt < 3$, contradicting the prescribed values. With non-strict convexity the data are possible: $f(t) = 2 + (t-2)/2 + (t-2)^2/6$ has affine derivative and satisfies all four endpoint values. The argument needs a strict loss of area, and the common endpoint upper bound of one in Figure 2 is only a convenience. The example is pure calculus, chosen to illustrate the chord argument.

With the right slope changed to 2, the chord has mean $5/4$, so a strict loss of area can leave the required mean one. Indeed,

$$f(t) = 2 + \frac{t-2}{2} + \frac{(t-2)^3}{18}$$

has the required endpoint values, slopes $1/2$ and 2, and $(f')'' = 1/3 > 0$. Thus these changed data are compatible. Strict convexity produces the contradiction only when the chord cannot supply enough area.

Question 4. Set $h = q' - 1$. Since $q(r) = r$, reaching or passing the diagonal at t means $\int_r^t h \geq 0$. Strict convexity places h below its endpoint chord, so that integral is less than $(t-r)(h(r)+h(t))/2$. With $h(r) \leq 0$, this forces $h(t) > 0$. The chord rule on $[r, s]$, solved for $h(s)$, keeps every later value at least $h(t)$. Integrating gives $q(s) \geq t + (s-t)(1 + h(t)) > 1$ for a suitable $s < 1$, violating the

bound. If q never reaches the diagonal, then $G(r) < 0$ instead. Thus both possibilities contradict the conditions on a minimiser.

For the proposed polynomial, $q(0) = 0$, $q'(0) = 1$, $q''' = 6 > 0$, and direct integration gives $G(s) = s^3(1-s)/4 \geq 0$, with $G(0) = 0$. Its derivative has minimum $13/16$, so q increases to $5/4$ as $t \uparrow 1$. It is bounded, yet $q(9/10) = 2043/2000 > 1$, so the other reduced conditions cannot replace the exact cap supplied by $|u| \leq 1$.

Question 5. Choose $r = 0.4$ and $\varepsilon = 0.02$. The theorem gives a pair differing on between 29% and 31% of sites. Reverse every sign of the second configuration. Its energy is unchanged, its overlap changes sign, and the differing-site fraction becomes its complement, between 69% and 71%. This operation uses the same successful instance and causes no further probability loss. The theorem fixes $\eta > 0$ before n grows and supplies no rate, so the stated result stops short of the stronger choice $\eta = 1/n$.

Further reading

Full support alone does not give a density. Additional differentiation of the identities $q = t$ and $\mathbb{E}[a^2] = 1$, followed by another use of the integrated curvature estimate on locally constant comparison coefficients, proves that γ is smooth, $\gamma(0) = 0$, and $\mu_\gamma = \gamma'(t)\,dt$ with $\gamma'(t) > 0$ for $0 < t < 1$ (OpenAI (2026a), Corollary 6.4). The *Long Reasoning Record*, §§9.1–9.2, supplies this additional argument. Chen's earlier theorem had already given smooth nonnegative density and full support.

Further martingale arguments recover the ground-state value from the terminal slope response. They use the first curvature moment $\mathbb{E}[a]$, distinct from the second moment $\mathbb{E}[a^2] = 1$ used here. See OpenAI (2026a), Corollary 7.1 and the *Long Reasoning Record*, §6.

The separate circuit theorem in OpenAI (2026b) resolves the eventual-optimality conjecture for the Quantum Approximate Optimization Algorithm (QAOA) in Basso–Farhi–Marwaha–Villalonga–Zhou, §7, equation (7.2). The *Long Reasoning Record*, §8, develops its circuit construction and the order of limits. It requires additional arguments beyond full support.

Formal verification and sources

These notes reconstruct selected parts of OpenAI's analytic manuscript, and the global scaling argument for contact follows Chen's variational proof. The empty-support and terminal convexity arguments of Section 5.4 are ordinary deductions from the stated analytic inputs, and the finite no-tie argument of Section 6 follows from the Gaussian Hamiltonian. These deductions are proved here on paper.

OpenAI's formalisation, written in Lean 4 (de Moura–Ullrich, 2021) with mathlib (The mathlib Community, 2020), contains a formal statement of the support and moment identities of Section 4 for a minimiser, and Lean checks OpenAI's proof of the corresponding theorem. The statement takes the minimiser as a hypothesis, so the external Parisi formula and its attainment lie outside it. A separate formal statement concerns the value identity and the finite Gaussian approximation, and its proof has not been checked. Neither statement covers positive density, the near-ground overlap theorem used in Section 6 or the complete QAOA circuit theorem. The *Long Reasoning Record*, §13, compares the formal statements with the printed ones and gives the versions used. The links below select the editions of the sources.

- **OpenAI (2026a).** *Full support of the zero-temperature Sherrington–Kirkpatrick order parameter*, 27 September 2026. Pinned manuscript. §2 supplies regularity, Proposition 5.1 the strict curvature estimate, §6 the support proof, and Corollary 6.4 the additional density conclusion.
- **Hong-Bin Chen (2026).** *FRSB in the SK spin glass: convergence to full-interval support at zero temperature*. Version 1, Theorem 1.3 for earlier full support and smooth nonnegative density; Proposition 3.2 for contact by scaling the coefficient. Strictly positive density in the further-reading note uses the separate OpenAI corollary.
- **Antonio Auffinger and Wei-Kuo Chen (2017).** *Parisi formula for the ground state energy in the mixed p-spin model*, Annals of Probability 45(6B), 4617–4631. Version 2. The variational formula and attainment used in Section 3.2.
- **Aukosh Jagannath and Ian Tobasco (2016).** *A dynamic programming approach to the Parisi functional*, Proceedings of the American Mathematical Society 144, 3135–3150. Version 4. The control method, with finite-temperature terminal data; it is not by itself the zero-temperature regularity theorem used here.
- **Wei-Kuo Chen, Madeline Handschy and Gilad Lerman (2018).** *On the energy landscape of the mixed even p-spin model*, Probability Theory and Related Fields 171, 53–95. Version 4. Proposition 4 gives the first-variation formula used in Section 5.1; the proof of Theorem 6 gives the zero-candidate descent.
- **Antonio Auffinger, Wei-Kuo Chen and Qiang Zeng (2020).** *The SK model is infinite step replica symmetry breaking at zero temperature*. Version 2, Proposition 1 and Remark 7. Near-ground overlap pairs and the earlier terminal-plateau observation.
- **OpenAI (2026b).** *QAOA attains the SK ground-state energy in the thermodynamic-first limit*, 25 September 2026. Pinned manuscript. The separate circuit theorem.
- **Joao Basso, Edward Farhi, Kunal Marwaha, Benjamin Villalonga and Leo Zhou (2022).** *The Quantum Approximate Optimization Algorithm at High Depth for MaxCut on Large-Girth Regular Graphs and the Sherrington-Kirkpatrick Model*. Full preprint, version 3, §7, equation (7.2). The eventual-optimality conjecture.
- **Yuxin Zhou (2026).** *On the Gardner Transition in the Ising Pure p-Spin Glass II*. Version 1, Appendix E.3, (E.37) and (E.41). Earlier inequalities for the density comparison.