First Proof: Comparison of Our Solutions with the Reference Solutions (Q1–Q10)

February 14, 2026

Scope and format

This document consolidates (in order Q1–Q10) my prior comparisons between:

- **our LaTeX solutions** (the ten "q#_proof*.tex" drafts), and
- the **reference solutions/comments** in `FirstProofSolutionsComments.pdf`.

For each question, I list:

1. what the **reference** solution does (key objects, lemmas, spine),
2. what **our** solution does (as written),
3. a **detailed differences / missing pieces** checklist.

1 Question 1 (Hairer): non-(quasi)invariance of the Φ^4_3 measure under shifts

Reference solution: what it does

- Works with the *stochastic quantization* stationary process $u(t)$ whose fixed-time law is μ.
- Uses the decomposition $u = -\Phi + v$ where Φ solves the linearized SPDE and v is controlled by a *paracontrolled-type* expansion, e.g.

$$v = -3((v - \Psi) \prec \Phi) + v^\sharp$$

(cf. their Proposition 4.1-style statement).

- Introduces a key diverging constant $c_{N,2}$ and proves a lower bound $c_{N,2} \gtrsim \log N$ (their Lemma 4.3).
- Defines a *separating event* B_γ (for $\gamma > 1/2$) of the form

$$B_\gamma = \left\{ \lim_{N\to\infty} (\log N)^{-\gamma} \langle H_3(P_N u; c_N) + 9c_{N,2} P_N u, \psi\rangle = 0 \right\}$$

and proves $\mu(B_\gamma) = 1$ (via Lemmas 4.2, 4.4, 4.5 plus product/commutator estimates).

- Shows that under $u \mapsto u + \psi$ the extra term $9(\log N)^{-\gamma} c_{N,2}\langle \psi_N, \psi\rangle$ diverges for $1/2 < \gamma < 1$, so $T^*_\psi \mu(B_\gamma) = 0$.
- Concludes *mutual singularity* $\mu \perp T^*_\psi \mu$ for any nonzero smooth shift ψ.

Our solution: what it does

- Uses a mollifier scale $\varepsilon_n = \exp(-e^n)$ and an observable

$$X_n = \varepsilon_n^{3/4}\Big(\langle :u^3:_\varepsilon, \psi\rangle + C_2(\varepsilon)\langle u_\varepsilon, \psi\rangle\Big),$$

  and claims $X_n \to 0$ in probability under μ but $X_n \to \infty$ under the shifted measure due to a diverging linear counterterm.

- Implements the correct *high-level* idea: build an event that has μ-probability 1 but has $T_\psi^*\mu$-probability 0.

Differences / what we were missing (detailed checklist)

1. **Canonical construction not built:** the reference works through the stationary SPDE and the decomposition $u = -\Phi + v$. Our writeup does not establish (or precisely cite) this representation and the needed control space for v.
2. **Wrong/unspecified regularization scheme:** reference uses P_N and tracks $(\log N)^{-\gamma}$ with $N \to \infty$; we used a special subsequence ε_n without justifying equivalence of schemes or scaling.
3. **Missing the exact separating event B_γ:** reference uses the specific combination $H_3(P_N u; c_N)+ 9c_{N,2}P_N u$ and proves convergence to 0 after scaling; we only asserted a comparable combination with an unspecified $C_2(\varepsilon)$.
4. **Coefficient and Wick/Hermite structure:** the constant 9 and the precise correction $9c_{N,2}P_N u$ come from an explicit Wick/Hermite expansion and renormalization bookkeeping. We did not derive these constants.
5. **Key lemmas missing:** we did not prove the analogues of:
   - Lemma 4.2: scaled Wick-cubic decay/convergence ($\gamma > 1/2$ threshold),
   - Lemma 4.4: convergence of polynomial expressions in the cutoff field Φ_N,
   - Lemma 4.5: convergence of renormalized quadratic/cross terms,
   - Lemma 4.3: logarithmic divergence $c_{N,2} \gtrsim \log N$.
6. **Parameter regime clarity:** reference needs $1/2 < \gamma < 1$ to both kill the cubic term and make the drift blow up; our exponent 3/4 was hard-coded and not justified as part of a general γ-scheme.

2 Question 2 (Nelson): Rankin–Selberg test vector, single W for all π

Reference solution: what it does

- Constructs a **single** Whittaker function $W \in \mathcal{W}(\Pi, \psi^{-1})$ that works **uniformly for all** representations π of fixed conductor.
- Uses the **Godement–Jacquet functional equation** and **Mellin inversion**, defining a carefully chosen test function β and its transform $\beta^\sharp$ adapted to π's γ-factor.
- Reduces the shifted Rankin–Selberg integral to an integral over a **congruence subgroup** $K_1(\mathfrak{q})$ and uses **newvector theory** to prove the integral is **nonzero for all** s.
- Links the twist parameter u_Q to the conductor by choosing Q generating $\mathfrak{q}^{-1}$.

Our solution: what it does

- Starts from a K_n-fixed vector in π and constructs W by prescribing compactly-supported Kirillov/Whittaker data so that the integrand is "constant" on its support, turning the integral into a volume.

Differences / what we were missing (detailed checklist)

1. **Fatal requirement mismatch:** our W depends on π (through the chosen K_n-fixed vector). The problem requires **one** W depending only on Π, valid for **all** π of a given conductor.
2. **Constant-on-support strategy fails in general:** already for $n = 1$ the integral becomes a Gauss sum and the integrand cannot be constant; in higher rank central character obstructions persist (as stressed in the reference comments).
3. **Missing the core analytic engine:** no Godement–Jacquet functional equation, no Mellin inversion, no $\beta/\beta^\sharp$ mechanism tied to γ-factors.
4. **Missing the newvector/congruence subgroup reduction:** reference reduces to $K_1(\mathfrak{q})$ and uses newvector theory to force nonvanishing for all s; we do not.
5. **Conductor linkage not used:** the choice of Q is tied to π's conductor in the reference; our proof treats Q as an external parameter without exploiting ε-factor structure.

3 Question 3 (Williams): Markov chain with stationary measure from $F^*_\mu(\cdot; q{=}1, t)$

Reference solution: what it does

- Constructs a **nontrivial, explicit** chain ("interpolation t-PushTASEP"-type) whose transitions are defined **probabilistically/combinatorially**, not by quoting F^* to define the kernel.
- Proves stationarity using **two-line queues** / (signed) queue combinatorics and weight identities matching the interpolation ASEP/Macdonald objects specialized at $q = 1$.
- Uses the "restricted" hypothesis (unique 0, no 1's) as a structural input in the queue/weight algebra.

Our solution: what it does

- Proposes an adjacent transposition chain with swap probabilities involving rational functions like $\frac{x_i - t x_{i+1}}{x_i - x_{i+1}}$.
- Claims stationarity by an asserted exchange ratio identity $F^*_{s_i\mu}/F^*_\mu$ matching the swap-rate ratio.

Differences / what we were missing (detailed checklist)

1. **Wrong chain:** the reference chain is not a simple adjacent-swap process; it is a global push process with additional structure.
2. **Hidden dependence on the polynomials:** while our transition rule does not explicitly mention F^*, our verification relies on an unproved identity about $F^*_{s_i\mu}/F^*_\mu$; this is essentially the failure mode the comments highlight for "Metropolis–Hastings-like" approaches.

3. **Missing the combinatorial proof framework:** no two-line queues, no signed queues, no weight-generating argument.

4. **Kernel validity not established:** positivity/stochasticity (probabilities in $[0,1]$ and row-sums 1) is not proven carefully for the parameter domain.

4 Question 4 (Garza Vargas–Srivastava–Stier): finite free Stam inequality

Reference solution: what it does

- Works on ordered roots and defines the root map $\Omega_{\boxplus_n}(\alpha,\beta)=\gamma$ via

$$p \boxplus_n q \;=\; \sum_{\pi\in S_n}\prod_{i=1}^{n}(x-\alpha_i-\beta_{\pi(i)}).$$

- Defines the **score vector** $J_n(\alpha)_i=\sum_{j\neq i}(\alpha_i-\alpha_j)^{-1}$ and $\Phi_n(p)=\|J_n(\alpha)\|^2$.
- Uses reverse heat flow on roots (Lemma 1.1-style) and the key identity (Obs. 2.1)

$$J_{\boxplus_n}(aJ_n(\alpha),bJ_n(\beta))=(a+b)J_n(\gamma).$$

- Proves the **Jacobian contraction** on a codimension-2 subspace V:

$$\|J_{\boxplus_n}(u,v)\|^2\le\|u\|^2+\|v\|^2$$

  (Prop. 2.1), via a Hessian identity + **hyperbolic polynomial convexity** (Bauschke et al.).
- Concludes $(a+b)^2\Phi(\gamma)\le a^2\Phi(\alpha)+b^2\Phi(\beta)$ and optimizes a,b to obtain Stam.

Our solution: what it does

- Uses an operator/differential viewpoint, introduces a semigroup, and asserts a Blachman/Stam-type inequality from a "spectral lower bound" lemma, without establishing the Jacobian/Hessian/hyperbolicity machinery.

Differences / what we were missing (detailed checklist)

1. **Missing the central geometric object:** we did not build $\Omega_{\boxplus_n}$ and its Jacobian $J_{\boxplus_n}$ as the main proof driver.

2. **Missing Observation 2.1** (the way scores transform through $\Omega_{\boxplus_n}$ under reverse heat flow), which is the bridge from dynamics to the inequality.

3. **Missing the hard step (Prop. 2.1):** we did not prove the contraction inequality. The reference proves it via:

   - a Hessian identity expressing $J_{\boxplus_n}J^*_{\boxplus_n}$ in terms of Hessians of Ω_i on V, and
   - PSD-ness of certain Hessian combinations coming from hyperbolic polynomial convexity.

4. **Our "spectral lower bound" is not a substitute** for the Hessian/PSD argument and leaves a gap.

5. **Multiplicity handling:** reference deals with multiple roots via perturbation/continuity; we did not execute this rigorously.

5 Question 5 (Hill–Lawson–Hill): O-slice connectivity via geometric fixed points

Reference solution: what it does

- Works with an incomplete transfer system O and admissible H-sets $T = \coprod H/K_i$ with $K_i \to H$.
- Defines $\tau^O_{\geq n}$ as the localizing subcategory generated by norms $N^T S^1$ with $|T| \geq n$, using $N^T S^1 \simeq S^{\mathbb{R}\cdot T}$.
- Introduces the **characteristic subgroup** $\chi_O(H)$ and proves the numerical characterization (Theorem 2.7):

$$E \in \tau^O_{\geq n} \quad \iff \quad \forall H \leq G,\ [H : \chi_O(H)] \cdot \mathrm{gconn}(E)(H) \geq n.$$

- Proves the forward direction by computing geometric fixed points of generators (Lemma 2.3-type orbit count bound).
- Proves the converse using isotropy separation and a slice computation for geometric Mackey functors (Lemmas 2.5–2.6).

Our solution: what it does

- States a criterion in terms of geometric fixed points $\Phi^H(X)$ being $\lambda_H(n)$-connective, with

$$\lambda_H(n) = \min\{|T/H| : T \in O(H),\ |T| \geq n\}.$$

- Argues via isotropy separation + "geometric part" reduction, with a compressed localizing-subcategory argument.

Differences / what we were missing (detailed checklist)

1. **Missing the characteristic subgroup formulation:** we did not define $\chi_O(H)$ nor state/prove the clean formula with $[H : \chi_O(H)]$.
2. **Missing the link $\lambda_H(n) = \lceil n/[H : \chi_O(H)] \rceil$:** to reconcile our λ_H with the reference, one must prove this numerical identity and translate it into the $[H : \chi_O(H)] \cdot \mathrm{gconn}$ condition.
3. **Forward lemma missing:** we did not reproduce the reference's precise orbit-count inequality (geometric fixed points of the generators).
4. **Converse mechanism missing:** we did not reproduce the isotropy separation criterion plus the key lemma computing slices of $\Sigma^k HM$ (geometric Mackey functors).
5. **Categorical infrastructure omitted:** truncation/cover functors and discreteness results are part of the reference proof's spine (even if not strictly required for the bare equivalence).

6 Question 6 (Spielman): existence of an ε-light set of size $\geq c\varepsilon n$

Reference solution: what it does

- Proves the statement with an explicit constant, e.g. $c = 1/42$, giving $|S| \geq \varepsilon n/42$.
- Uses a **greedy iterative selection** with two maintained invariants:
  - a leverage-score (mass) control for the selected set,

  - a **spectral barrier potential** controlling the top eigenvalues throughout the process.

- Works in a normalized image space (e.g. Laplacian image), using eigenvalue inequalities (Ky Fan trace-type) and counting to show many "good" vertices exist at each step.

Our solution: what it does

- Proves an **upper bound obstruction**: any universal constant must satisfy $c \leq 1/2$ (perfect matching example).
- Gives a linearization reduction and isolates a conjectural subset-selection condition; shows an analogous statement fails for arbitrary PSD families.

Differences / what we were missing (detailed checklist)

1. **We do not prove any positive constant $c > 0$.** The reference proves existence with an explicit positive constant; our writeup only provides an obstruction and partial reductions.
2. **Missing the full constructive method:** no greedy algorithm with explicit invariants, no barrier potential, no iteration analysis establishing progress.
3. **Mismatch in technical route:** we relax to a linearized PSD-sum statement; the reference controls the actual L_S behavior via sharper spectral arguments tailored to the graph Laplacian structure.

7 Question 7 (Weinberger): uniform lattice with 2-torsion vs rationally acyclic universal cover

Reference solution: what it does

- Uses **lattice-specific structure**: reduces to an extension $\Gamma = \pi \rtimes \mathbb{Z}_2$ with π a torsion-free lattice.
- Employs **rigidity/higher-signature** technology (Novikov/assembly injectivity, symmetric signatures, and cobordism/fixed-set arguments) to contradict rational acyclicity of $\widetilde{M}$ under the presence of $\mathbb{Z}/2$-torsion in Γ.
- Key point: the obstruction is **not** purely Smith theory; it is a lattice/rigidity phenomenon.

Our solution: what it does

- Proves a *different* claim under a stronger hypothesis: if $\widetilde{M}$ is **integrally acyclic** (hence mod-p acyclic), then $\pi_1(M)$ is torsion-free, using Smith fixed-point theory plus a deck transformation argument.

Differences / what we were missing (detailed checklist)

1. **Hypothesis mismatch (fatal):** the problem assumes only $\mathbb{Q}$-acyclicity, not $\mathbb{Z}$-acyclicity. From $\mathbb{Q}$-acyclic you *cannot* deduce mod-2 acyclicity; thus Smith theory does not apply as we used it.
2. **We ignored the essential lattice input:** in the rational setting, the reference needs the lattice/semisimple structure and the rigidity tools; our proof does not use them.

3. **Known counter-phenomena:** the comments stress that many "pure Smith" arguments fail because there exist rationally acyclic spaces admitting free $\mathbb{Z}/2$-actions in other contexts; the lattice assumption is precisely what blocks these.

4. **Missing the higher-signature / assembly-map argument** that drives the reference proof.

8 Question 8 (Abouzaid): Lagrangian smoothing of a polyhedral Lagrangian surface

Reference solution: what it does

- Proves a sharp **linear-algebra normal form** at a vertex meeting four faces: after a linear symplectic change, the local model is a product of a standard "positive axes union" with $\mathbb{R}^2$ (Lemma 1-type).
- Deduces existence of a Lagrangian plane L so that the symplectic-pairing projection $\Sigma \to L^\vee$ is a **homeomorphism** (Corollary 1-type), not just transverse.
- Introduces **smoothing functions** $S(\Sigma)$ via a canonical C^1 function q_Σ (piecewise quadratic), and proves a bijection: smoothing functions (mod constants) $\leftrightarrow$ graphical Lagrangians near Σ (Lemma 3 / Lemma 6-type).
- Handles edges using a **contractibility** statement for choices of L and compatible local data (Lemma 4/5-type).
- Globalizes by constructing a **dual conormal fibration** L_z over all $z \in K$ (Definitions 3–4, Lemma 8-type).
- Proves existence of smoothing functions of **arbitrarily small C^1 norm** (Lemmas 7 and 9-type), ensuring the resulting Lagrangians remain graphical in the chosen neighborhood globally.
- Assembles a **Hamiltonian isotopy** by concatenating graphical Hamiltonian paths between successive smoothings (final assembly).

Our solution: what it does

- Uses a Maslov-cycle/transversality argument to get a common cotangent chart transverse to the four tangent planes.
- Writes the local model as a graph of df for a continuous piecewise quadratic f and mollifies f.
- Attempts global gluing by summing locally supported Hamiltonians from vertex/edge charts.

Differences / what we were missing (detailed checklist)

1. **Missing the stronger vertex normal form:** transversality of planes is not enough; the reference produces a projection that is a *homeomorphism* on the local quadrant-union model.

2. **No conormal fibration:** we did not construct the global family of fibers L_z; without it, "graphical over varying fibers" is not justified.

3. **Missing smoothing-function formalism:** we did not define q_Σ, $S(\Sigma)$, nor prove the bijection between smoothing functions and Lagrangian graphs.

4. **Global gluing gap:** "summing Hamiltonians" does not guarantee the intended local behavior on overlaps (flows do not commute). The reference avoids overlap issues by producing globally-defined smoothings with small C^1 norm.

5. **Missing the small C^1 norm existence argument:** essential to keep fibers disjoint and remain in the graphical neighborhood globally.

6. **Edge compatibility mechanism missing:** the reference proves contractibility/compatibility of choices along edges and collars; our proof only sketches local crease smoothing.

9 Question 9 (Miao–Lerman–Kileel): equations for scaled determinant tensors $Q^{(\alpha\beta\gamma\delta)}$

Reference solution: what it does

- Packages all blocks into a tensor $Q \in \mathbb{R}^{3n\times 3n\times 3n\times 3n}$.
- Proves a **Tucker decomposition**

$$Q = C \times_1 A \times_2 A \times_3 A \times_4 A,$$

  where $A = [A^{(1)}; \ldots ; A^{(n)}] \in \mathbb{R}^{3n\times 4}$ and $C \in \mathbb{R}^{4\times 4\times 4\times 4}$ is a universal "sign/permutation" core (Lemma 1-type). Concludes multilinear rank $\leq (4, 4, 4, 4)$.

- Defines F to be **exactly the set of all** 5×5 **minors** of the four mode-flattenings (degree 5, independent of n).
- "If" direction: rank-1 scaling off-diagonal corresponds to Tucker scaling by invertible diagonal matrices, preserving multilinear rank, hence all minors vanish.
- "Only if" direction: assumes all minors vanish, normalizes λ, and uses **three explicit minor computations** (Steps 1–3) to force a rigid pattern:
  - entries with two "1"-indices are a constant c,
  - with one "1"-index are c^2,
  - with no "1"-indices are c^3,

  yielding $\lambda = u \otimes v \otimes w \otimes x$ off-diagonal.

Our solution: what it does

- Defines F as *swap quadrics* (Plücker-style) **plus** the 5×5 flattening minors.
- Proves the forward direction by exhibiting a rank-≤ 4 factorization for one flattening (Hodge-star/inner-product argument).
- For the reverse direction, attempts a tangent-space/Hadamard-stabilizer approach and only claims identifiability on the fully distinct-index set $\mathcal{I}_{\mathrm{obs}}$.

Differences / what we were missing (detailed checklist)

1. **Definition of F differs:** reference uses *only* the flattening minors. Our "swap quadrics" are not used/needed and complicate the system.

2. **Missing the full Tucker decomposition (all modes):** we sketched a mode-wise rank argument, but did not present the universal core C and the simultaneous multilinear-rank bound $(4, 4, 4, 4)$.

3. **Scope mismatch on indices (major):** the theorem concerns *all off-diagonal* tuples (not all equal), including repeated-index patterns. Our reverse direction treats only pairwise distinct indices $\mathcal{I}_{\text{obs}}$, which is strictly weaker.
4. **Missing the minors-to-λ recovery:** the reference uses explicit 5×5 determinants to force algebraic equalities among λ entries (Steps 1–3). Our tangent-space argument does not yield these global equalities.
5. **Unproven stabilizer lemma:** our reverse direction relies on an asserted classification of Hadamard deformations, which is not established and not the reference method.

10 Question 10 (Kolda): PCG for RKHS tensor mode update (matrix-free MVP + preconditioning)

Reference solution: what it does

- Notes the symmetric system may be **indefinite** and adds a ρI regularization to enforce SPD.
- Uses $K = UDU^T$ and transforms to a system in $\bar{W}$.
- Defines restricted objects $\hat{Z}$ and (row-wise) Kronecker structure so each row of F is a **row-wise Kronecker product**.
- Proves key lemmas enabling fast operations:
  - fast computation of Cx exploiting row-wise Kronecker structure,
  - formula for $C^T v$,
  - fast computation of $\text{diag}(C^T C)$.
- Uses a **diagonal preconditioner** based on $\text{diag}(\bar{F}^T \bar{F}) + \lambda(I \otimes D) + \rho I$.
- Gives explicit complexity: per-iteration MVP $O(qnr)$ and storage $O(q(n + r))$.

Our solution: what it does

- Derives the matrix-free gather/scatter MVP for $F^T F x$ and the RKHS regularization term.
- Proposes a Kronecker-structured preconditioner by approximating $\mathcal{P}_\Omega \approx \alpha I$ and using eigendecompositions of K and $Z^T Z$.

Differences / what we were missing (detailed checklist)

1. **Definiteness treatment:** reference explicitly adds $\rho > 0$ to guarantee SPD even when K is PSD/singular or the system is otherwise indefinite; we did not treat this carefully.
2. **Preconditioner choice and justification:**
   - reference: cheap *diagonal* preconditioner from explicit $\text{diag}(C^T C)$ lemmas,
   - ours: heavier Kronecker eigen-based preconditioner, not the one justified in the reference.
3. **Missing the explicit lemma suite:** we used the idea but did not state/prove the clean algebraic formulas for Cx, $C^T v$, and $\text{diag}(C^T C)$ that the reference highlights.
4. **Avoiding big-M costs:** we suggested forming $Z^T Z$ with complexity depending on the full product size M; the reference emphasizes staying in observed-index structures and avoiding large intermediate constructs.