The Sign Problem in Quantum Monte Carlo

Xiao Yan Xu

July 10, 2026

Abstract

Quantum Monte Carlo methods turn quantum many-body partition functions into statistical sums, but the resulting configuration weights need not be non-negative. This article reviews how negative and complex weights arise in stochastic-series-expansion, world-line, and determinant formulations; how modulus and real-part reference ensembles quantify their cancellations; and which structural principles can remove or mitigate the obstruction. The average sign conventionally decays exponentially as a free-energy difference, but special algebraic cases have only polynomial sign-induced cost. The emphasis is on representation dependence: the sign problem is not a property of the spectrum alone, but of the basis, auxiliary-field channel, contour, or resummation used to sample it. The article also summarizes complexity-theoretic limits and scoped topological obstructions that show why no universal classical cure should be expected.

Contents

1 Introduction

Quantum Monte Carlo (QMC) is powerful when a quantum partition function can be written as a sum or integral over configurations with non-negative weights. The fermion sign problem is the failure of precisely that condition. Its most useful formulation is not as a mysterious feature of fermions, but as a statement about the chosen representation.

Central point.

The sign problem is a property of the ***representation****, not of the physics.* Its severity is quantified by the average sign

$$\langle \mathrm{sgn} \rangle \;=\; e^{-\beta N\,\Delta f_N}, \tag{1}$$

where β is the inverse temperature, N the system size, and $\Delta f_N \geq 0$ a representation-dependent finite-size free-energy-density difference. If Δf_N approaches a positive constant, reweighting has an exponential cost; if it vanishes as $(\log N)/(\beta N)$, the sign instead decays algebraically. Positive representations are sometimes available by design; generic efficient curing is hard, and specified local classes can be excluded by phase invariants.

Equation (1) is the spine of the analysis below. When the thermodynamic free-energy difference is nonzero, the number of Monte Carlo samples needed at fixed relative accuracy grows exponentially in the space–time volume. Much of modern algorithm design searches for a representation in which that difference vanishes or is parametrically smaller.

Two organizing questions.

(Q1) **Where do negative weights come from?** (their microscopic origin: fermionic exchange, frustration, discretization, contour choice).

(Q2) **When — and how hard is it — to make them go away?** (design principles, symmetry criteria, and complexity-theoretic obstructions).

Representative settings. Two distinctions are useful but not exhaustive. First, a real sign $W \in \mathbb{R}$ requires only a $\mathbb{Z}_2$ reweighting factor, whereas a complex weight $W = |W|e^{i\theta}$ has a continuous phase. Second, the sampled degrees of freedom may be spins/bosons or fermions. This produces the examples in Fig. 1: frustrated spin SSE and real determinant signs in the left column; Berry/real-time phases and complex determinants in the right. Real time occurs for bosons and fermions. Likewise, finite chemical potential does not automatically imply a complex determinant: doped Hubbard DQMC commonly has real signed weights, and special finite-density gauge theories retain conjugation-based positivity. The classification belongs to the chosen representation, not to the Hamiltonian alone. Pedagogical overviews of the sign problem and many-body QMC are given by Pan and Meng [1], Xu [2], Assaad and Evertz [3]; sign-free fermion design is reviewed by Li and Yao [4]. Complex actions and phase-mitigation methods are surveyed by Alexandru et al. [5], Berger et al. [6], Aarts and Sexty [7].

sampled field	**character of the sampled weight** real sign	 complex phase
spin/boson	frustrated spin/boson SSE $W \in \mathbb{R}$, $W \lessgtr 0$	Berry or real-time phases $W = |W|e^{i\theta}$
fermion	real fermion determinants doped Hubbard is a common case	complex fermion determinants when conjugation protection is lost

Figure 1: Representative, representation-level sources of real and complex weights. The boxes are examples rather than exclusive physical classes: real-time evolution spans both rows, while finite density can land in either fermionic column depending on its antiunitary and conjugation symmetries.

Notation and conventions

We work in natural units where $\hbar = k_B = 1$ throughout, so temperature is an energy and $\beta = 1/T$; the physical constants are restored only where a factor is genuinely at stake. A Monte Carlo *configuration* $\mathcal{C}$ is a point in whatever sampling space the chosen representation defines (an auxiliary-field history, a world-line configuration, an operator string). The partition function is written as a sum over configurations of a weight $W(\mathcal{C})$, and the sign is factored out as $\mathrm{sgn}(\mathcal{C}) = \mathrm{sgn}\, W(\mathcal{C})$. Table 1 collects the recurring symbols.

Symbol	Meaning
β	inverse temperature, $\beta = 1/T$ (with $k_B = 1$)
N	system size (number of sites / spatial volume)
$\Delta\tau$	imaginary-time (Trotter) step
$L_\tau = \beta/\Delta\tau$	number of imaginary-time (Trotter) slices
Z	partition function, $Z = \mathrm{Tr}\; e^{-\beta H}$
$W(\mathcal{C})$	weight of configuration $\mathcal{C}$
$\mathrm{sgn}(\mathcal{C}) = \mathrm{sgn}\, W(\mathcal{C})$	sign of that weight
$\langle \mathrm{sgn} \rangle$	average sign of the sampled ensemble
$Z_{\lvert W \rvert}$, $Z_{\lvert \mathrm{Re}\, W \rvert}$	positive references built from $\lvert W \rvert$ and, when Z is real, $\lvert \mathrm{Re}\, W \rvert$
Δf_N	finite-size free-energy-density difference between Z and the chosen positive reference

Table 1: Notation used throughout the article.

Section 2.3 derives reweighting and its statistical cost. Before doing so, one distinction will prevent several common overstatements: finding non-negative weights removes the *sign-induced* variance, but does not by itself prove that a Markov chain mixes in polynomial time. We therefore use *sign-free representation* for $W(\mathcal{C}) \geq 0$ and reserve *polynomial-time cure* for an estimator whose total cost, including autocorrelations and the cost of constructing the representation, is polynomial. *Mitigation* means retaining a residual sign/phase or accepting a controlled approximation.

2 Origin of the Sign Problem and the Stoquastic Criterion

2.1 Stochastic series expansion and the birth of the sign

We now watch a sign problem appear *ex nihilo*, in the cleanest bosonic setting available: the stochastic series expansion (SSE) representation of a lattice spin Hamiltonian. SSE is a finite-temperature quantum Monte Carlo method built directly on the high-temperature series of the partition function, without any Trotter discretization [8, 9]. It is the ideal minimal laboratory: the configuration weight is a bare product of Hamiltonian matrix elements, so the positivity requirement is completely transparent, and the sign problem, when it strikes, is impossible to hide behind formalism.

2.1.1 The Taylor expansion of the partition function

Start from the object every finite-T method must estimate,

$$Z = \operatorname{Tr} e^{-\beta H} = \sum_{n=0}^{\infty} \frac{(-\beta)^n}{n!} \operatorname{Tr}\big[H^n\big] = \sum_{n=0}^{\infty} \frac{\beta^n}{n!} \operatorname{Tr}\big[(-H)^n\big]. \tag{2}$$

The last rewriting is deliberate: it is $-H$, not H, whose matrix elements will play the role of Boltzmann weights, and the reader should already suspect that the fate of the method hinges on the *signs* of the matrix elements of $-H$.

To make (2) sampleable we do two things: decompose H into a sum of elementary *Hamiltonian bond operators*, and insert resolutions of the identity in a convenient local basis.

Bond decomposition. Write

$$H = \sum_b H_b, \tag{3}$$

where each H_b acts nontrivially only on the few sites belonging to a bond (or plaquette) b. Expanding $(-H)^n$ using this bond decomposition turns the n-th term of (2) into a sum over ordered *bond strings*

$$S_n = (b_1, b_2, \dots, b_n), \tag{4}$$

each string being a particular way of choosing which Hamiltonian bond term H_b appears at each of the n propagation steps. The operator that appears in the SSE weight is explicitly $-H_b$.

Complete sets of states. Choose a basis $\{|\alpha\rangle\}$ of *local product states* — for spins, the Ising basis $|\alpha\rangle = |\sigma_1^z \sigma_2^z \cdots \sigma_N^z\rangle$. Inserting $\mathbb{1} = \sum_\alpha |\alpha\rangle \langle\alpha|$ between every pair of adjacent operators, and using the cyclic trace to close the chain, we obtain

$$\operatorname{Tr}\big[(-H)^n\big] = \sum_{S_n} \sum_{\{\alpha_0,\dots,\alpha_n\}} \Big(\prod_{p=1}^{n} \langle\alpha_p|(-H_{b_p})|\alpha_{p-1}\rangle \Big) \delta_{\alpha_n,\alpha_0}. \tag{5}$$

The Kronecker delta $\delta_{\alpha_n,\alpha_0}$ is the fingerprint of the trace: the sequence of *propagated states*

$$|\alpha_0\rangle \xrightarrow{-H_{b_1}} |\alpha_1\rangle \xrightarrow{-H_{b_2}} \cdots \xrightarrow{-H_{b_n}} |\alpha_n\rangle = |\alpha_0\rangle \tag{6}$$

must be *periodic*. A full SSE configuration is thus the pair $(S_n, \{\alpha\})$: an operator string together with the state it propagates around the imaginary-time loop.

2.1.2 The configuration weight

Collecting the prefactor from (2), each configuration carries the weight

$$\boxed{W(S_n, \{\alpha\}) = \frac{\beta^n}{n!} \prod_{p=1}^{n} \langle\alpha_p|(-H_{b_p})|\alpha_{p-1}\rangle} \tag{7}$$

so that $Z = \sum_{S_n,\{\alpha\}} W$. Monte Carlo sampling requires that we interpret W/Z as a probability distribution over configurations, and this is legitimate *only if* $W \geq 0$ *for every configuration.* The combinatorial factor $\beta^n/n!$ is manifestly positive, so positivity of the weight is entirely controlled by the product of matrix elements.

Two kinds of factors appear. A *diagonal* element $\langle\alpha|(-H_b)|\alpha\rangle$ can always be made non-negative by subtracting a constant from H (equivalently, shifting $-H \to -H + C$ and absorbing that constant into the bond operators), which merely rescales Z by $e^{\beta C}$ and never changes an expectation value. The dangerous factors are the *off-diagonal* ones, $\langle\alpha'|(-H_b)|\alpha\rangle$ with $\alpha' \neq \alpha$, which flip the local state and cannot be shifted away. Their signs depend on the chosen basis and decomposition, but not on a diagonal energy offset.

Termwise SSE positivity criterion. A simple sufficient condition for every weight in (7) to be non-negative is that every off-diagonal matrix element of each $-H_b$ be non-negative,

$$\langle\alpha'|(-H_b)|\alpha\rangle \geq 0 \qquad \text{for all } \alpha' \neq \alpha, \tag{8}$$

equivalently, every off-diagonal matrix element of each H_b is non-positive. This *termwise stoquastic* condition is stronger than necessary: selection rules can force negative factors to occur in pairs, or several operator strings can be regrouped into a positive contribution. It is nevertheless the standard local certificate because it makes positivity manifest configuration by configuration.

Remark 2.1. The criterion is basis dependent. Marshall rotations search for a better product basis; auxiliary-field, Majorana, and fermion-bag methods change the configuration space more radically. Searching specified families of local stoquastic bases can be NP-hard, but it is only one kind of cure.

2.1.3 Worked example: a single Heisenberg bond

Take the antiferromagnetic Heisenberg interaction on one bond $\langle ij \rangle$,

$$H_b = J\,\mathbf{S}_i \cdot \mathbf{S}_j = J\Big[S_i^z S_j^z + \tfrac{1}{2}(S_i^+ S_j^- + S_i^- S_j^+)\Big], \qquad J > 0, \tag{9}$$

with spin-$\frac{1}{2}$ operators $\mathbf{S} = \frac{1}{2}\boldsymbol{\sigma}$ (we keep the factor of $\frac{1}{2}$ explicit; the ladder operators satisfy $S^\pm = S^x \pm iS^y$). In the local S^z basis $\{|\uparrow\uparrow\rangle, |\uparrow\downarrow\rangle, |\downarrow\uparrow\rangle, |\downarrow\downarrow\rangle\}$ the only nonzero off-diagonal action is the spin exchange,

$$\langle\uparrow\downarrow|H_b|\downarrow\uparrow\rangle = \frac{J}{2}\,\langle\uparrow\downarrow|S_i^+ S_j^-|\downarrow\uparrow\rangle = \frac{J}{2}, \qquad \langle\downarrow\uparrow|H_b|\uparrow\downarrow\rangle = \frac{J}{2}. \tag{10}$$

For antiferromagnetic $J > 0$ this off-diagonal element is *positive*, hence the corresponding element of $-H_b$ is $-J/2 < 0$: the bond is not termwise stoquastic as written. This does *not* mean that an isolated bond already has a negative closed string—periodicity forces its exchange operator to occur an even number of times. The issue is whether one basis choice fixes all bonds simultaneously on an extended graph.

The Marshall/sublattice rotation. On a *bipartite* lattice ($A \cup B$, every bond connecting A to B) we perform the Marshall sign transformation [10], a π rotation about the z-axis of every spin on sublattice B,

$$S_j^x \to -S_j^x, \quad S_j^y \to -S_j^y, \quad S_j^z \to S_j^z \qquad (j \in B), \tag{11}$$

which is the canonical (unitary) transformation $U = \prod_{j\in B} e^{i\pi S_j^z}$. Because $S^\pm = S^x \pm iS^y$, this sends $S_j^\pm \to -S_j^\pm$ for $j \in B$, so the transverse term acquires an overall minus sign,

$$\tfrac{1}{2}(S_i^+ S_j^- + S_i^- S_j^+) \;\longrightarrow\; -\tfrac{1}{2}(S_i^+ S_j^- + S_i^- S_j^+), \tag{12}$$

while the diagonal $S_i^z S_j^z$ is untouched. The rotated bond off-diagonal element becomes

$$\langle\uparrow\downarrow|\tilde{H}_b|\downarrow\uparrow\rangle = -\frac{J}{2} < 0 \quad\Longrightarrow\quad \langle\uparrow\downarrow|(-\tilde{H}_b)|\downarrow\uparrow\rangle = +\frac{J}{2} \geq 0. \tag{13}$$

The "negative-of-a-positive" has become the positive we needed. Every bond of a bipartite Heisenberg antiferromagnet is simultaneously curable by the *same* global rotation, so SSE for the square-, cubic-, or honeycomb-lattice Heisenberg model is rigorously sign-free — a positivity that underlies the loop/operator-loop update [8, 11] and, more broadly, the "Marshall-positive" family of $SU(N)$ magnets that can be simulated with controlled statistics [12, 13].

2.1.4 Contrast: the frustrated triangle

The cure above relied crucially on bipartiteness: the sublattice rotation exists because the lattice two-colors. Take instead three antiferromagnetically coupled spins on a *triangle* (the elementary plaquette of the triangular lattice). Two-color the vertices however you like — say $1 \in A$, $2 \in B$; then the third vertex is forced to be adjacent to a same-color neighbor no matter how it is labeled. One of the three bonds always connects two spins of the same sublattice, and the Marshall rotation that fixes the other two bonds *fails* on it, leaving a residual positive off-diagonal element and hence sign-indefinite operator strings (Fig. 2). This is the geometric origin of the sign problem for frustrated quantum magnets in the site basis. It rules out the Marshall two-coloring, not every possible reformulation. For a broad family of frustrated $S = \frac{1}{2}$ antiferromagnets one can trade the site basis for the basis of spin eigenstates of small clusters, in which the competing exchanges no longer generate positive off-diagonal elements and much of the sign problem simply evaporates [14, 15]. Special meron constructions provide another exact but model-dependent route [16]. The quantitative cost of a residual sign is derived in Sec. 2.3.

Proposition 2.1 (Sufficient condition for SSE sign-freeness)**.** If there exists a local unitary $U = \prod_i u_i$ (a product of on-site rotations, generally sublattice dependent) such that every bond operator of the transformed Hamiltonian $\tilde{H} = UHU^\dagger$ satisfies $\langle\alpha'|\tilde{H}_b|\alpha\rangle \leq 0$ for all $\alpha' \neq \alpha$ in the product basis, then all SSE weights (7) are non-negative and the simulation is free of the sign problem. For the nearest-neighbor Heisenberg antiferromagnet this U is the Marshall rotation and exists precisely on bipartite lattices.

The triangle shows how a negative *closed product* arises in a chosen SSE decomposition and why the Marshall gauge cannot repair it. It does not establish that every local basis or regrouping fails; that stronger quantifier belongs to the later complexity and intrinsic-obstruction discussions.

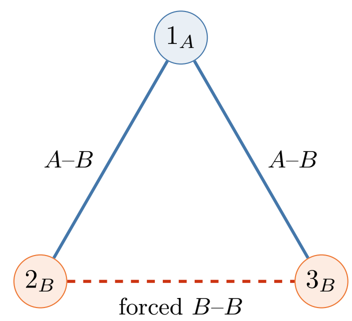

solid: $\langle\alpha'|-\widetilde{H}_b|\alpha\rangle = +J/2$ dashed: $\langle\alpha'|-\widetilde{H}_b|\alpha\rangle = -J/2$

Figure 2: The frustrated triangular plaquette of antiferromagnetically coupled spins. A sublattice (Marshall) rotation with $1 \in A$, $2 \in B$ renders the two A–B bonds termwise stoquastic (solid), but the third bond is forced to join equal colors. Its positive off-diagonal element survives (dashed), showing that no Marshall $\mathbb{Z}_2$ assignment works on an odd cycle. The figure does not rule out non-diagonal or cluster-basis reformulations.

2.2 What stoquasticity does—and does not—guarantee

The SSE example motivates a representation-level definition that will be used throughout the article.

Definition 2.1 (Stoquastic Hamiltonian)**.** A Hamiltonian H is *stoquastic* in a specified product basis $\{|\alpha\rangle\}$ when

$$\langle\alpha|H|\alpha'\rangle \in \mathbb{R}, \qquad \langle\alpha|H|\alpha'\rangle \leq 0 \quad (\alpha \neq \alpha'). \tag{14}$$

A local decomposition is *termwise stoquastic* when every H_b has this property in the same basis.

After shifting diagonal terms by constants, termwise stoquasticity makes every factor in Eq. (7) non-negative. Stoquasticity of the full matrix H also implies that $C\mathbb{1} - H$ is entrywise non-negative for a large enough C. Perron–Frobenius then guarantees, within each irreducible block, a lowest-energy state whose amplitudes can be chosen non-negative,

$$|\psi_0\rangle = \sum_\alpha c_\alpha\,|\alpha\rangle\,, \qquad c_\alpha \geq 0. \tag{15}$$

Disconnected blocks may remain degenerate. Thus positivity is compatible with topological sectors: for example, the toric code is manifestly stoquastic in the standard basis, while local operators do not connect its distinct noncontractible-loop sectors [17].

Two limitations are essential. First, non-negative weights remove cancellation but do not guarantee rapid mixing; autocorrelation times or other conditioning problems can still be exponential. Polynomial mixing has to be established separately (it is known for some operator-loop algorithms [18]). Second, stoquasticity is sufficient, not necessary. Regrouping operator strings, changing the auxiliary-field algebra, or resumming fermions can produce positive weights even when the written Hamiltonian is not stoquastic. Twisted quantum doubles, discussed in Sec. 3.7, provide a particularly sharp example. Accordingly, finding a stoquastic basis is one route to a sign-free representation, not its definition.

2.3 Reweighting, average signs, and cost

Let $Z = \sum_{\mathcal{C}} W(\mathcal{C}) > 0$, as for the thermal partition function of a Hermitian Hamiltonian. If W is real, the usual positive reference ensemble is

$$Z_{|W|} = \sum_{\mathcal{C}} |W(\mathcal{C})|, \qquad p_{|W|}(\mathcal{C}) = \frac{|W(\mathcal{C})|}{Z_{|W|}}. \tag{16}$$

Writing $W = s|W|$, with $s = \operatorname{sgn} W$, gives

$$\langle \mathrm{sgn} \rangle \equiv \langle s \rangle_{|W|} = \frac{Z}{Z_{|W|}}, \qquad 0 < \langle \mathrm{sgn} \rangle \le 1, \tag{17}$$

and the exact reweighting identity

$$\boxed{\langle O \rangle = \frac{\langle O\, s \rangle_{|W|}}{\langle s \rangle_{|W|}}.} \tag{18}$$

Complex weights with real Z. The standard phase-quenched construction samples $Z_{|W|} = \sum_{\mathcal{C}} |W|$ and reweights by $e^{i\theta} = W/|W|$. However, the known reality of Z permits a better reference:

$$\begin{aligned} Z &= \sum_{\mathcal{C}} \operatorname{Re} W(\mathcal{C}), \\ Z_{|\operatorname{Re} W|} &= \sum_{\mathcal{C}} |\operatorname{Re} W(\mathcal{C})| = \sum_{\mathcal{C}} |W(\mathcal{C})|\, |\cos\theta(\mathcal{C})|, \\ s_R(\mathcal{C}) &= \operatorname{sgn}\big(\operatorname{Re} W(\mathcal{C})\big). \end{aligned} \tag{19}$$

Sampling $p_R = |\operatorname{Re} W|/Z_{|\operatorname{Re} W|}$ gives the *sign-quenched* average $\mathcal{A}_R = \langle s_R \rangle_R = Z/Z_{|\operatorname{Re} W|}$ [19]. By contrast, the phase-quenched average is $\mathcal{A}_{\mathrm{pq}} = \langle e^{i\theta} \rangle_{|W|} = \langle \cos\theta \rangle_{|W|} = Z/Z_{|W|}$; its imaginary part vanishes because Z is real.

Pointwise, $|\operatorname{Re} W| \leq |W|$. Therefore

$$Z \leq Z_{|\operatorname{Re} W|} \leq Z_{|W|}, \qquad 0 < \mathcal{A}_{\mathrm{pq}} \leq \mathcal{A}_R \leq 1, \tag{20}$$

and, more explicitly,

$$\mathcal{A}_{\mathrm{pq}} = \mathcal{A}_R \left\langle |\cos\theta| \right\rangle_{|W|}. \tag{21}$$

Thus the conventional average phase is a lower bound on the real-part average sign. Whenever the reality of Z can be exploited, $|\operatorname{Re} W|$ has the larger reweighting signal.

Often the reality is manifest through a measure-preserving involution $\mathcal{C} \mapsto \bar{\mathcal{C}}$ with $W(\bar{\mathcal{C}}) = W(\mathcal{C})^*$; a conjugate pair then contributes $2 \operatorname{Re} W$. For a real, conjugation-even estimator, Eq. (18) holds with $W \to \operatorname{Re} W$, $s \to s_R$, and the R ensemble. More generally one must pair the numerator as $2 \operatorname{Re}[WO]$, rather than blindly replacing W by $\operatorname{Re} W$ in every observable. The construction is unavailable for a genuinely complex target, and zeros of $\operatorname{Re} W$ can create sampling barriers. A larger average sign also need not imply better mixing or smaller variance for every observable.

Free-energy and cost scaling. Choose $Z_X = Z_{|W|}$ or $Z_X = Z_{|\operatorname{Re} W|}$, and write $Z_X = e^{-\beta N f_X}$. Then

$$\mathcal{A}_X \equiv \frac{Z}{Z_X} = e^{-\beta N \Delta f_{X,N}}, \qquad \Delta f_{X,N} \equiv f - f_X \geq 0. \tag{22}$$

Equation (20) is equivalently $0 \leq \Delta f_{R,N} \leq \Delta f_{\mathrm{pq},N}$. The free-energy identity is exact, but exponential decay is only the conventional case: $\Delta f_{X,N} \to \Delta f_X > 0$ gives $\mathcal{A}_X \sim e^{-\beta N \Delta f_X}$. In special models,

$$\mathcal{A}_X \sim N^{-\alpha} \quad \Longleftrightarrow \quad \Delta f_{X,N} \sim \frac{\alpha \log N}{\beta N}, \tag{23}$$

so the sign penalty is algebraic rather than exponential. Such behavior occurs in projected extended-Hubbard and moiré-motivated models; sign-bound theory also identifies conditions under which an algebraic lower bound is possible [1, 20, 21].

For a unit-modulus sign or phase factor r, let $\widehat{\mathcal{A}}_X = N_s^{-1} \sum_{i=1}^{N_s} r_i$. Independent samples satisfy

$$\mathbb{E}\Big[|r - \mathcal{A}_X|^2\Big] = 1 - \mathcal{A}_X^2, \qquad \mathbb{E}\Big[|\widehat{\mathcal{A}}_X - \mathcal{A}_X|^2\Big] = \frac{1 - \mathcal{A}_X^2}{N_s}. \tag{24}$$

With correlated Markov-chain samples one replaces N_s by $N_{\text{eff}} \simeq N_s/(2\tau_{\text{int}})$. Hence

$$\frac{\delta \mathcal{A}_X}{\mathcal{A}_X} \simeq \frac{e^{\beta N \Delta f_{X,N}}}{\sqrt{N_{\text{eff}}}} \quad (\mathcal{A}_X \ll 1), \tag{25}$$

and fixed relative precision ϵ requires

$$\boxed{N_{\text{eff}} \gtrsim \frac{1}{\epsilon^2 \mathcal{A}_X^2} = \frac{1}{\epsilon^2} e^{2\beta N \Delta f_{X,N}}.} \tag{26}$$

For the algebraic case $\mathcal{A}_X \sim N^{-\alpha}$, this becomes $N_{\text{eff}} \gtrsim \epsilon^{-2} N^{2\alpha}$: polynomial sign-induced cost, provided configuration evaluation and mixing are also polynomial. The numerator of Eq. (18) is correlated with the denominator, so the prefactor is observable dependent; barring a special cancellation, the same inverse-$\mathcal{A}_X$ scaling remains. Slow mixing increases the raw number of updates by the additional factor $2\tau_{\text{int}}$. This separates two bottlenecks that are often conflated: sign cancellations reduce N_{eff}, while poor mixing reduces how fast independent samples are produced.

Figure 3 displays the relation on logarithmic axes. The intersection with the estimator's standard error, $N_{\text{eff}}^{-1/2}$, occurs at $x_\star = \frac{1}{2} \log N_{\text{eff}}$, where $x = \beta N \Delta f_{X,N}$.

The free-energy difference depends on both the representation and its positive reference. A basis change, auxiliary-field channel, analytic pairing, resummation, or contour deformation can change $\Delta f_{X,N}$, sometimes to zero. Conversely, the measured sign can carry physical information: its finite-size scaling may locate phase transitions or obey useful bounds relative to a sign-free reference system [21, 22]. These uses do not remove the sampling penalty, but they explain why the sign itself is sometimes a meaningful observable rather than only a diagnostic.

2.4 The second face: fermion determinants

In the preceding SSE discussion, signs arose from products of Hamiltonian matrix elements. For interacting fermions the same cancellations can be resummed into determinants, moving the sign rather than eliminating it.

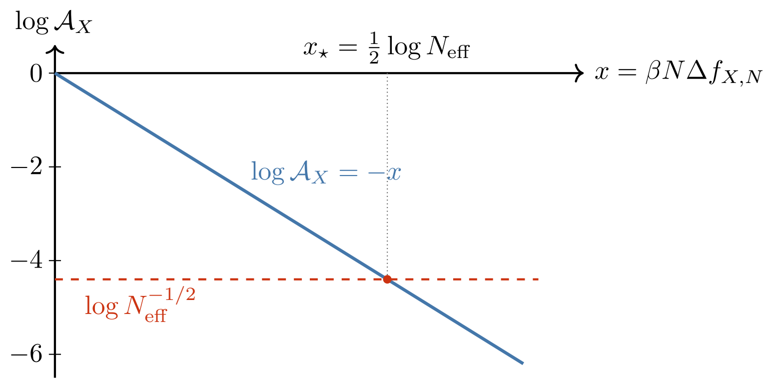

Figure 3: The reweighting wall in representation-neutral variables. The per-configuration sign or phase has order-one variance, but the standard error of its mean is $N_{\text{eff}}^{-1/2}$. Once $\mathcal{A}_X = e^{-x} \lesssim N_{\text{eff}}^{-1/2}$, the denominator is unresolved. Exponential decay with system size is the conventional case $x \propto \beta N$; algebraic decay instead has $x \sim \alpha \log N$. The horizontal line is an estimator error, not single-sample noise.

2.4.1 Why fermions are generically non-stoquastic

For fermions the obstruction is structural. Write any second-quantized operator in an occupation-number basis $|n_1 n_2 \cdots\rangle$ with a fixed site ordering. A hopping term $c_i^\dagger c_j$ acting on such a state produces not merely a relabeled occupation but a Jordan–Wigner string:

$$c_i^\dagger c_j \, |\ldots n_i \ldots n_j \ldots\rangle = (-1)^{\sum_{i<k<j} n_k} \, |\ldots (n_i{+}1) \ldots (n_j{-}1) \ldots\rangle \,. \tag{27}$$

The prefactor $(-1)^{\sum n_k}$ is the parity accumulated by dragging the annihilation operator past the fermions sitting between sites j and i. In the worldline picture, a closed spacetime configuration in which particles have permuted their identities carries a global factor $(-1)^P$, where P is the parity of that permutation. Two worldlines that exchange contribute -1; this is Fermi statistics made graphical.

Remark 2.2 (An open-chain exception). For the nearest-neighbor Hubbard chain with open boundaries, order the modes consecutively within each spin species. Every same-spin hop then joins adjacent modes, and same-species worldlines can neither exchange nor wind around a boundary. The occupation/worldline representation is therefore sign-free at any filling and for either sign of U [3]. A ring, longer-range hopping, or a ladder removes this no-exchange proof; it does not follow that every auxiliary-field determinant representation of the open chain is configuration-wise positive.

A minimal fermionic loop. The obstruction already appears for two spinless fermions on a three-site ring. In the ordered basis $(|12\rangle, |13\rangle, |23\rangle)$, with $|ij\rangle = c_i^\dagger c_j^\dagger |0\rangle$ for $i < j$, uniform microscopic hopping $-t$ induces

$$H_t = t \begin{pmatrix} 0 & -1 & +1 \\ -1 & 0 & -1 \\ +1 & -1 & 0 \end{pmatrix}. \tag{28}$$

The configuration-graph loop product is $(-t)(-t)(+t) = +t^3$. It is invariant under any diagonal rephasing because the three phases telescope, whereas three strictly negative off-diagonal elements would have product $-t^3$. Thus no diagonal phase gauge makes this occupation-basis matrix stoquastic. This compact example rules out that specific cure, not a non-diagonal basis or a regrouped expansion.

In higher-dimensional and multiply connected hopping graphs, such exchange loops are generic rather than exceptional. Fermionic Hamiltonians are therefore generically non-stoquastic in occupation bases, although the severity and even the existence of a sign problem remain representation dependent [23, 24].

The standard escape is not to fight the permutation signs term by term but to *resum* them. An antisymmetrized sum over single-particle labelings is precisely a determinant,

$$\sum_{P \in S_N} (-1)^P \prod_k \langle \phi_k \,|\, \psi_{P(k)} \rangle = \det\big[\langle \phi_i | \psi_j \rangle\big], \tag{29}$$

the Leibniz formula for the determinant. Determinant QMC (DQMC) is built on this identity: it trades an intractable, sign-alternating sum over $N!$ worldline permutations for a single $N \times N$ determinant that can be evaluated in polynomial time. The signs have not vanished — they now live in the *sign of the determinant*. That is the second face.

2.4.2 Auxiliary-field / determinant QMC in one pass

We illustrate on the Hubbard model,

$$H = -t \sum_{\langle ij \rangle, \sigma} (c_{i\sigma}^\dagger c_{j\sigma} + \text{h.c.}) + U \sum_i n_{i\uparrow} n_{i\downarrow}, \tag{30}$$

with $\sigma = \uparrow, \downarrow$ and $n_{i\sigma} = c_{i\sigma}^\dagger c_{i\sigma}$. Before decoupling, the quartic interaction prevents the fermionic trace from being Gaussian. DQMC [25] Trotter-factorizes the propagator, introduces an auxiliary field that makes the fermion action bilinear, integrates out the fermions, and samples the remaining field. The same construction underlies ground-state projector AFQMC [26, 27].

Trotter split. Write $e^{-\beta H} = \left(e^{-\Delta\tau H}\right)^{L_\tau}$ with $\beta = L_\tau\,\Delta\tau$, and separate hopping K from interaction V,

$$e^{-\Delta\tau H} = e^{-\Delta\tau K/2} e^{-\Delta\tau V} e^{-\Delta\tau K/2} + O(\Delta\tau^3). \tag{31}$$

The symmetric factorization gives a global $O(\Delta\tau^2)$ error at fixed β; convergence in $\Delta\tau \to 0$ must still be checked.

Discrete Hubbard–Stratonovich (Hirsch). The interaction on each site is decoupled by an *Ising* auxiliary field $s_i = \pm 1$ rather than a continuous field [28]. Shifting to the particle–hole symmetric form $U\,n_\uparrow n_\downarrow \to U(n_\uparrow - \frac{1}{2})(n_\downarrow - \frac{1}{2})$ (plus a chemical-potential shift), the key exact identity for the repulsive case is

$$e^{-\Delta\tau U\,(n_\uparrow-\frac{1}{2})(n_\downarrow-\frac{1}{2})} = \frac{1}{2}\,e^{-\Delta\tau U/4} \sum_{s=\pm1} e^{\lambda s\,(n_\uparrow - n_\downarrow)}, \qquad \cosh\lambda = e^{\Delta\tau U/2}. \tag{32}$$

Since $n_\sigma \in \{0,1\}$, the identity follows by matching the empty/doubly occupied and singly occupied sectors. The field couples to the *spin* density for repulsive $U > 0$; for attractive $U < 0$ a real density-channel decoupling couples it to $n_\uparrow + n_\downarrow - 1$. Choosing an exact HS channel is a change of auxiliary-field representation, not a gauge transformation. It leaves Z invariant while changing the configuration-wise sign structure [29].

For repulsive $U > 0$ there is also a *density-channel* decoupling, but it is complex:

$$e^{-\Delta\tau U\,(n_\uparrow-\frac{1}{2})(n_\downarrow-\frac{1}{2})} = \frac{1}{2}\,e^{+\Delta\tau U/4} \sum_{s=\pm1} e^{i\gamma s\,(n_\uparrow + n_\downarrow - 1)}, \qquad \cos\gamma = e^{-\Delta\tau U/2}. \tag{33}$$

Before any particle–hole transformation, both spins couple identically to the imaginary auxiliary potential $i\gamma s$, and the factor $e^{-i\gamma s}$ is an imaginary configuration-dependent constant. This complex density channel is the cleanest way to see the half-filled bipartite conjugate-pair proof below. These discrete HS formulas are exact operator identities: the four local occupation states reduce to two sectors that can be matched directly. The DQMC discretization error comes from the Trotter split, not from the HS transformation.

Integrating out the fermions. After (32) the exponent is bilinear in $c^\dagger, c$ for every fixed configuration $\{s_{i,\ell}\}$ (site i, time slice ℓ). Each spin species propagates independently through a product of one-body matrices,

$$B_\ell^\sigma[\{s\}] = e^{-\Delta\tau K/2}\, e^{V_\ell^\sigma[\{s\}]} e^{-\Delta\tau K/2}, \qquad M_\sigma[\{s\}] = \mathbb{1} + B_{L_\tau}^\sigma B_{L_\tau-1}^\sigma \cdots B_1^\sigma. \tag{34}$$

The Gaussian (Grassmann) integral over the fermions yields a determinant for each species, and the partition function becomes a classical sum over Ising configurations [30, 31]:

$$\begin{aligned} Z &= \sum_{\{s\}} W[\{s\}], \\ W[\{s\}] &= \det M_\uparrow[\{s\}] \det M_\downarrow[\{s\}]. \end{aligned} \tag{35}$$

The Monte Carlo weight is W. The sign problem is now sharp and concrete: W *can be negative*, and the second face is nothing but $\det M_\uparrow \det M_\downarrow < 0$. For a systematic, modern account of this world-line/determinantal machinery see Assaad and Evertz [3].

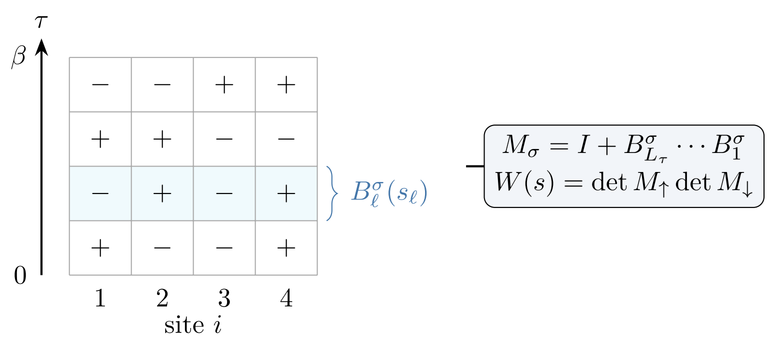

Figure 4: DQMC auxiliary-field geometry. Every cell carries an Ising field $s_{i,\ell} = \pm 1$; one highlighted time slice determines the one-body propagator B_ℓ^σ. The ordered product over slices forms M_σ, and integrating out the two spin species gives the sampled determinant weight.

2.4.3 Sign-free cases to memorize

Three benchmark cases organize single-band Hubbard DQMC. The first two have configuration-wise positivity proofs; the third lacks that protection.

(i) Attractive $U < 0$, any filling, spin balanced. Decoupling in the *density* (charge) channel makes the auxiliary field couple to $n_\uparrow + n_\downarrow$ identically for both spins. Provided the two species also have identical one-body terms (in particular, a common chemical potential and no Zeeman imbalance), $M_\uparrow[\{s\}] = M_\downarrow[\{s\}]$. Hence

$$W[\{s\}] = \big(\det M_\uparrow[\{s\}]\big)^2 \;\geq\; 0 \qquad \text{for all } \{s\}. \tag{36}$$

No particle–hole symmetry is required: the spin-balanced attractive Hubbard model is sign-free at arbitrary total density and temperature. Spin-dependent chemical potentials destroy the square argument.

(ii) Repulsive $U > 0$, half filling, bipartite lattice. There are two standard positivity arguments. First, a particle–hole transformation on one spin species maps the half-filled repulsive Hamiltonian to an attractive Hubbard Hamiltonian (up to harmless constants), so case (i) applies. Second, start from the complex density-channel HS form (33). The field initially couples identically to the two spins and carries the imaginary constant $e^{-i\gamma s}$; after the bipartite particle–hole transformation on the down spin, $n_{i\downarrow} \to 1 - n_{i\downarrow}$, this constant cancels and the two spins couple with opposite imaginary potentials, $+i\gamma s$ and $-i\gamma s$. Thus the spin determinants are complex conjugates configuration by configuration:

$$\det M_\downarrow[\{s\}] = \big(\det M_\uparrow[\{s\}]\big)^* \implies W = \big|\det M_\uparrow\big|^2 \;\geq\; 0. \tag{37}$$

(The general antiunitary pairing argument is given in Sec. 3.2.3.)

(iii) Repulsive $U > 0$, doped away from half filling. Neither configuration-wise pairing argument applies. The conventional DQMC representation is therefore generically sign-problematic (though special parameters and finite systems can be exceptions). This is the regime relevant to doped-Mott, pseudogap, and high-T_c physics.

How the average sign degrades away from half filling depends strongly on temperature, interaction, lattice geometry, and the chosen representation; it is not a universal proxy for correlation strength [32].

2.4.4 One phenomenon, two representations

In an SSE/worldline representation the pathology is a negative closed product of matrix elements; in a determinantal representation it is a negative or complex determinant. By the Leibniz identity (29), a determinant is itself a signed sum over permutation histories. The two descriptions are related by resummation, although their configuration spaces and average signs need not coincide.

One phenomenon. SSE signs and negative or complex determinants are different representations of one obstruction: $W(\mathcal{C})/Z$ is not a non-negative real probability measure. Consequently:

- A basis, HS channel, or resummation changes the configurations and can change the sign problem even though it leaves physical observables invariant.
- In any *fixed* representation, its reweighting cost is governed by that representation and reference's $\Delta f_{X,N}$ in Eq. (22). There is no representation-independent average sign: a successful reformulation changes $\Delta f_{X,N}$ or makes it vanish.

Once a representation is fixed, the corresponding sign-, real-part-, or phase-reweighting identity applies regardless of whether the obstruction came from a permutation, a frustrated operator string, or a determinant. The remaining sections ask how a better representation can avoid or reduce the wall, and what limits such a search.

3 Design Principles and Fundamental Limits

The rest of the article separates three logically different tasks: constructing an exact positive representation, mitigating a residual sign or phase, and establishing limits on what any prescribed class of reformulations can achieve. Keeping these tasks distinct is essential: failure to find a cure is not a no-go theorem, and positivity alone is not a proof of polynomial runtime.

3.1 Taxonomy of responses

A cure need not be a local basis rotation. It may instead use determinant pairing, a cluster cancellation, or an exact resummation. Conversely, a method that stabilizes a calculation by imposing a trial-state constraint is an approximation, not a cure. Table 2 records the distinctions used below.

The sections that follow first collect constructive positivity certificates, then compare exact cancellations with controlled approximations and contour methods, and finally state the complexity and phase-level no-go results with their scopes.

Response	Goal, examples, and limitation
Exact positive representation	Reexpress Z with $W(\mathcal{C}) \geq 0$, using (for example) Marshall rotations, determinant/Kramers pairing, fermion bags, Majorana or split-orthogonal positivity, or exact meron cancellation. This removes sign cancellation; mixing and construction cost remain separate.
Mitigation / approximation	Retain a residual phase or replace it by trial-state bias. Reweighting, finite-flow or learned contours, complex Langevin, fixed-node, constrained-path, and phaseless methods belong here. Exponential variance or systematic error may remain.
Limits and obstructions	Establish worst-case hardness or exclude positive representatives within a stated local class. NP-hard signed instances, hard basis searches, and phase invariants have different scopes; failure to find a cure is not a certificate.

Table 2: Three distinct responses to a sign or phase problem.

3.2 Design principles for sign-free models

We now collect structural certificates that make a chosen representation positive configuration by configuration. They are sufficient, not necessary: a model outside every class listed here may still admit a different positive expansion.

3.2.1 Positivity versus reality: sign problem vs. phase problem

Before cataloguing principles it is worth making a distinction that is often blurred in informal discussions. In determinant QMC the weight of an auxiliary-field configuration ϕ is a product of fermion determinants, $W(\phi) = \det M_\uparrow(\phi) \det M_\downarrow(\phi)$, up to a positive bosonic factor. Two qualitatively different pathologies can occur:

- **Sign problem:** the weights are *real* but not sign-definite, $W(\phi) \in \mathbb{R}$, $W(\phi) \lesseqgtr 0$. One samples $|W|$ and reweights by $\operatorname{sgn} W = \pm 1$.
- **Phase problem:** the weights are genuinely complex. For a complex target one samples $|W|$ and retains the full $U(1)$ phase. When the target partition function is known to be real, conjugate symmetrization instead permits the sign-quenched measure $|\operatorname{Re} W|$. This converts the continuous phase into a residual sign with the hierarchy in Eq. (20), but does not by itself remove exponential cancellations.

The two are governed by different algebraic requirements, and it pays to keep them separate:

> *Reality* of a determinant follows from the *weaker* statement that the spectrum of the relevant (generally non-Hermitian) fermion matrix is invariant *as a set* under complex conjugation, $\lambda \mapsto \lambda^*$, so that $\det M = \prod_i \lambda_i \in \mathbb{R}$. This still permits *real* eigenvalues, which may be negative — hence reality alone does not fix the sign. *Positivity* requires additional structure, such as Kramers pairing, a Pfaffian square, or the connected-component theorem for the split-orthogonal group. Reciprocal eigenvalue pairing by itself is not enough: for example, $(1-2)(1-1/2) < 0$. Reality removes a phase problem; positivity removes the residual sign problem.

Configuration-wise determinant reality is stronger than merely knowing that the total Z is real. The latter permits the conjugate-pair regrouping of Eq. (19), which can still leave both positive and negative real parts. We will see the algebraic dichotomy concretely in §3.2.6: the *split orthogonal* group forces $\det M \geq 0$ (kills the sign problem outright), whereas a *pseudo-unitary* constraint only forces $\det M \in \mathbb{R}$ (kills the phase problem, leaving a possible residual sign).

3.2.2 Marshall and Hubbard pairing

The Marshall rotation of Sec. 2.1 is the canonical bosonic certificate: on a bipartite nearest-neighbor antiferromagnet it changes every transverse matrix element from $+J/2$ to $-J/2$ in one common product basis. Its $SU(N)$ extensions support broad families of sign-free loop models [12]. Odd cycles obstruct this diagonal two-coloring, but not necessarily cluster bases or regrouped expansions.

For the single-band Hubbard model, the corresponding DQMC certificate is configuration-wise determinant pairing. The two cases established above are

$$M_\uparrow(s) = M_\downarrow(s) \quad \Longrightarrow \quad W(s) = [\det M_\uparrow(s)]^2 \geq 0, \tag{38}$$

$$\det M_\downarrow(s) = [\det M_\uparrow(s)]^* \quad \Longrightarrow \quad W(s) = |\det M_\uparrow(s)|^2 \geq 0. \tag{39}$$

The first applies to the spin-balanced attractive model at arbitrary filling; the second follows at half filling for repulsive U on a bipartite lattice after the appropriate particle–hole transformation. Both statements require the relation to hold for every auxiliary-field configuration, not merely after averaging.

Lieb's spin-reflection positivity is historically related but should not be used as the microscopic DQMC proof of these identities. It proves ground-state spin and uniqueness results for Hubbard-type Hamiltonians [33, 34]; DQMC needs the stronger configuration-wise matrix relation displayed above.

3.2.3 Time-reversal / Kramers positivity (Wu–Zhang)

Let $B(\phi) = \prod_\ell e^{-\Delta\tau h(\phi_\ell)}$ and $M(\phi) = I + B(\phi)$. Suppose the *same* antiunitary operator $T = U_T K$ satisfies, for every time slice and every fixed auxiliary-field configuration,

$$T h(\phi_\ell) T^{-1} = h(\phi_\ell), \qquad T^2 = -I. \tag{40}$$

Then nonreal eigenvalues μ of B occur with their conjugates, while real eigenvalues have even algebraic multiplicity. Consequently

$$\det M(\phi) = \prod_a |1 + \mu_a|^2 \geq 0. \tag{41}$$

The proof is the Kramers argument in one-particle space. If $Bv = \mu v$, antilinearity and $TB = BT$ give $B(Tv) = \mu^*(Tv)$. For real μ, the vectors v and Tv are independent, since $Tv = cv$ would imply $T^2 v = |c|^2 v$, contradicting $T^2 = -I$. The same map pairs generalized eigenspaces, so the conclusion also holds for a nondiagonalizable B. This is the Wu–Zhang condition [35].

The configuration-wise qualifier is crucial. A physical time-reversal symmetry of the interacting Hamiltonian is useful only if the chosen HS decomposition preserves it for each ϕ. A common chemical potential is time-reversal even and does not by itself spoil the criterion; Zeeman fields or HS channels that distinguish a Kramers pair do. The construction underlies sign-free metallic quantum-critical models [36–38] and does not require S^z conservation.

3.2.4 Fermion bags: resummation before sampling

The half-filled spinless t–V model on a bipartite lattice has only one complex-fermion species, so no spin determinant can square the weight:

$$H = -t \sum_{\langle ij \rangle} (c_i^\dagger c_j + \text{h.c.}) + V \sum_{\langle ij \rangle} (n_i - \tfrac{1}{2})(n_j - \tfrac{1}{2}). \tag{42}$$

The fermion-bag construction instead performs a selected part of the fermionic sum exactly *before* Monte Carlo sampling [39, 40]. In the continuous-time expansion, a k-bond configuration produces a $2k \times 2k$ matrix

$$G([s]) = D_0([s]) + A, \qquad (D_0)_{qq'} = \frac{s_q}{2} \delta_{qq'},$$

where $s_q = \pm 1$ labels particle/hole insertions and A contains off-diagonal free propagators. Summing the 2^{2k} labels analytically gives

$$\sum_{[s]} \det[D_0([s]) + A] = 4^k \det A. \tag{43}$$

Those local diagonal blocks are the bags in this formulation.

Bipartiteness supplies the remaining sign. With $\widetilde{D}_{qq'} = \sigma_{i_q} \delta_{qq'}$, where $\sigma_i = \pm 1$ on the two sublattices, $A^T = -\widetilde{D} A \widetilde{D}$. Hence $A\widetilde{D}$ is real antisymmetric and

$$(-1)^k \det A = \det\left(A\widetilde{D}\right) = \mathrm{Pf}(A\widetilde{D})^2 \geq 0. \tag{44}$$

Together with the expansion coefficient, this makes every resummed weight non-negative for $V > 0$. The result enabled sign-free simulations of the honeycomb spinless critical point [41].

More generally a bag is any subset of Grassmann variables or worldline segments summed exactly so that the remaining weights factorize positively; it need not be a literal connected spatial region. This viewpoint reaches odd-flavour and impurity models beyond ordinary determinant pairing [42–44].

3.2.5 Majorana representation and Majorana positivity

The next viewpoint on the same spinless-fermion obstruction was to recast the problem in *Majorana* rather than complex-fermion variables. Writing $c_j = \frac{1}{2}(\gamma_j^1 + i\gamma_j^2)$ with Hermitian Majoranas $\gamma^2 = 1$, Li et al. [45] introduced Majorana QMC and showed that the half-filled spinless t–V family can also be made sign-free by decoupling the interaction in a Majorana channel: the two Majorana species play the role of conjugate partners, so the weight becomes a non-negative square. This was historically after the fermion-bag solution, but conceptually important because it turned the spinless/odd-species cure into a determinant-style auxiliary-field principle.

The underlying algebra was then organized group-theoretically by Li et al. [46] into the *Majorana-time-reversal* classes: within that symmetry classification there are two fundamental sign-free classes — the *Kramers class*, positive by the $T^2 = -1$ degeneracy of §3.2.3, and a genuinely new *Majorana class*, positive by decoupling the Majorana bilinear into a time-reversal-partner block pair B and $\bar{B}$ (needing neither charge conservation nor a single $T^2 = -1$), which reaches spin–orbit-coupled and some doped models.

Wei et al. [47] sharpened this into the notion of *Majorana positivity*: a sufficient condition on the coefficient matrices of the Majorana Hamiltonian guaranteeing $W(\phi) \geq 0$ that unifies essentially *all previously known* sign-free auxiliary-field models (the authors note it does not exhaust all positive decompositions). It supplies two distinct criteria — a Majorana *reflection* positivity (a positive-semidefinite-kernel argument) and a Majorana *Kramers* positivity (an antiunitary $T = SC$ with $T^2 = -1$) — and carves out a strictly larger sign-free territory than the complex-fermion principles alone. For relativistic (Dirac) actions, the Majorana-positivity condition has also been recast as a Pfaffian semi-positivity criterion and used to classify sign-free cases such as QCD with chiral or isospin chemical potential [48]. The algebraic backbone was later cast as a *contraction semigroup* of imaginary-time propagators — a maximal unitary subgroup together with an invariant positivity cone — which unifies the Kramers, Majorana, split-orthogonal, and fermion-bag criteria as special cases [49].

3.2.6 Group-theoretic guiding principles

The Majorana and semigroup viewpoints suggest a general strategy: identify a *Lie group* $\mathcal{G}$ such that all imaginary-time propagators $e^{-\Delta\tau\, h(\phi_l)}$ lie in $\mathcal{G}$, and let the representation theory of $\mathcal{G}$ dictate the analytic structure of $\det M$. Two such groups make the positivity/reality distinction of §3.2.1 razor-sharp.

Split orthogonal group → positivity. Let $B(\phi)$ denote the ordered one-body propagator. If every real generator obeys $\eta A \eta = -A^T$, with $\eta = \mathrm{diag}(I_n, -I_n)$, then the product lies in the identity component $O_{++}(n,n)$. The connected-component theorem gives

$$B(\phi) \in O_{++}(n,n) \quad \Longrightarrow \quad \det[I + B(\phi)] \geq 0. \tag{45}$$

This is stronger than reciprocal eigenvalue pairing: the orientation of the identity component controls negative-real-axis blocks even when the product is not one exponential [50]. It explains the positivity of the half-filled spinless t–V model without Kramers degeneracy.

Pseudo-unitary group → reality only. If instead $B^\dagger \eta B = \eta$ and $\det B = 1$, as in a determinant-one pseudo-unitary group, then

$$\begin{aligned} \overline{\det(I+B)} &= \det\Big(I + B^\dagger\Big) = \det\Big(I + \eta B^{-1} \eta^{-1}\Big) \\ &= \det\Big(B^{-1}\Big) \det(I+B) = \det(I+B). \end{aligned} \tag{46}$$

Thus $\det(I+B) \in \mathbb{R}$, but it can be negative. Pseudo-unitarity removes a phase problem; the $O_{++}(n,n)$ orientation theorem supplies the additional positivity in the split-orthogonal case [51].

3.2.7 Summary of positivity principles

The principles above are best remembered as a dictionary from a structural input (a symmetry or group constraint) to the analytic guarantee it buys, together with the models where it is deployed.

Positivity toolbox.

Structural input	Guarantee and representative use
Stoquastic / Marshall basis	$\langle\alpha\| e^{-\beta H} \|\alpha'\rangle \geq 0$; bipartite $SU(N)$ antiferromagnets.
Hubbard determinant pairing	$M_\uparrow = M_\downarrow$ or conjugate determinants; attractive and half-filled bipartite repulsive Hubbard models.
Kramers $T^2 = -1$	$\det M \geq 0$; Dirac and quantum-critical fermions without requiring S^z conservation.
Fermion bags / Majorana positivity	Positive resummed determinants or Pfaffian squares; spinless and odd-flavour models, including some spin–orbit-coupled cases.
Split orthogonal $O_{++}(n,n)$	$\det M \geq 0$; spinless t–V and group-designed models.
Pseudo-unitary $\mathrm{SU}(m,n)$	$\det M \in \mathbb{R}$ only; removes a phase but can leave a sign, as in QED_3 parent-state constructions.

3.2.8 Outlook: from sufficient conditions to bounds

All of the above are binary guarantees: a chosen formulation either satisfies a positivity criterion or it does not. Sign bounds provide quantitative information in between. For a determinant D averaged with a positive base measure P, the Hermitian-system sign is

$$S_R = \frac{\langle \mathrm{Re}\, D\rangle_P}{\langle |\, \mathrm{Re}\, D|\rangle_P}.$$

Because $\langle |\, \mathrm{Re}\, D|\rangle_P \leq \langle |D|\rangle_P$,

$$S_R \geq \frac{\langle \mathrm{Re}\, D\rangle_P}{\langle |D|\rangle_P} = \frac{Z_D}{Z_{|D|}}. \tag{47}$$

This is the determinant version of Eq. (20): the $|D|$ ratio is a lower bound on the better $|\,\mathrm{Re}\, D|$ average sign. When $Z_{|D|}$ is the partition function of a physical sign-free reference Hamiltonian, its low-temperature behavior is controlled by

$$\frac{Z_D}{Z_{|D|}} \sim \frac{g_D}{g_{|D|}}\, e^{-\beta(E_D - E_{|D|})},$$

where E and g are ground-state energies and degeneracies [21]. An algebraic lower bound excludes an exponentially small S_R; an exponentially small lower bound is inconclusive rather than evidence that the actual sign must be exponential.

Positivity symmetries can also constrain the physics accessible to the sign-free model. Antiunitary protection imposes correlation-function constraints and can force a symmetry-distinguished operator to become critical at a transition [52]. Within every presently known DQMC sign-free class, a stable Fermi liquid in $d \geq 2$ is impossible unless the protecting antiunitary symmetry is spontaneously broken [53]. These are constraints on specified positivity structures, analogous in spirit to the scoped topological obstructions of Sec. 3.7.

3.3 Exact cancellations and trial-state constraints

Meron clusters and constrained-path methods are often listed together because both stabilize signful simulations, but they have different logical status. A meron construction analytically cancels signed partners and is exact when its factorization conditions hold. Fixed-node, constrained-path, and phaseless methods restrict the sampled paths using a trial state and therefore introduce a systematic approximation.

3.3.1 The meron-cluster algorithm

Cluster algorithms decompose a configuration into independently flippable loops or clusters [8, 11]. In a meron representation, choose a reference orientation with sign sgn_0 and let $f_c = 0, 1$ record whether cluster c is flipped. The sign factorizes as

$$\mathrm{sgn}(\mathcal{C}(\{f_c\})) = \mathrm{sgn}_0 \prod_c \sigma_c^{f_c}, \qquad \sigma_c \in \{+1, -1\}, \tag{48}$$

where σ_c is the sign change under flipping cluster c [54].

Definition 3.1 (Meron)**.** A cluster c whose flip *reverses* the sign of the configuration, i.e. one that carries $\sigma_c = -1$ so that flipping it sends $\mathrm{sgn}(\mathcal{C}) \to -\mathrm{sgn}(\mathcal{C})$, is called a *meron.* A cluster whose flip leaves the sign invariant ($\sigma_c = +1$) is a non-meron.

The useful cases have three additional properties: every cluster can be flipped independently without changing $|W|$, a reference orientation has non-negative weight, and the improved estimator is compatible with the observable. The analytic flip average is

$$\sum_{\{f_c\}} W(\{f_c\}) = |W|\ \mathrm{sgn}_0 \prod_{c=1}^{n_c} \sum_{f_c=0,1} \sigma_c^{f_c} = |W|\ \mathrm{sgn}_0 \prod_{c=1}^{n_c} (1 + \sigma_c). \tag{49}$$

Any meron supplies a factor $1 + \sigma_c = 0$, so the entire flip orbit cancels from the partition function.

When independent cluster flips preserve $|W|$, the sign factorizes as in Eq. (48), and the zero-meron reference weight is non-negative, every nonzero-meron orbit cancels in the flip-averaged partition function. The surviving zero-meron sector is then an exact positive representation. Observable numerators may require additional sectors, commonly zero and two merons, with their own improved estimators.

One simple physical example is the semi-frustrated quantum spin model treated by Henelius and Sandvik [16]. In an S^z basis, take an XXZ-type Hamiltonian

$$H = \sum_{i,r} \left[J_z(r) S_i^z S_{i+r}^z + \frac{J_{xy}(r)}{2} (S_i^+ S_{i+r}^- + S_i^- S_{i+r}^+) \right],$$
$$J_z(r) < 0, \qquad J_{xy}(r) = -J_z(r) > 0. \tag{50}$$

The positive transverse matrix elements produce a sign problem in the usual SSE/operator-loop formulation. The loop decomposition, however, makes the sign flip under a loop update visible: a loop whose flip changes the configuration sign is a meron. For the partition function and observables such as the internal energy, all nonzero-meron sectors cancel by (49), so the simulation can be restricted to the positive zero-meron sector. More general observables may require sampling zero- and two-meron sectors; away from the exact meron point, merons can still reduce the sign problem without eliminating its exponential scaling [55].

The exponential gain is decisive: instead of averaging a signal $\langle \mathrm{sgn} \rangle$ that vanishes exponentially, one samples a positive measure directly. The price is *model specificity*. The construction requires (i) a cluster representation in which the sign factorizes as in (48), and (ii) that the meron cancellation aligns with the observable of interest. This holds for particular lattice fermion models and quantum spin systems — e.g. staggered fermions in a (3+1)D chiral phase transition [56], certain t–V and quantum-antiferromagnet models — but is *not* a universal cure. When it applies, however, it is *exact*: there is no systematic bias. A related, also-exact and also-model-specific cure is the *fermion-bag* resummation of §3.2.4, which localizes and cancels the sign within compact space-time regions and solves sign problems for a class of models unreachable by the traditional determinantal criteria [42].

3.3.2 Constrained-path and fixed-node methods

When no sign-free representation is known, one may instead accept a systematic approximation. The oldest incarnation is the *fixed-node* method in continuum and lattice diffusion Monte Carlo, built on fermion nodes [57]. The antisymmetric ground state $\Psi_0(\mathbf{R})$ changes sign across the nodal surface $\{\mathbf{R} : \Psi_0(\mathbf{R}) = 0\}$. If that surface were known, projection within one nodal pocket would be exact. In practice one imposes the nodes of a trial wavefunction Ψ_T.

Theorem 3.1 (Fixed-node variational upper bound; ten Haaf et al.)**.** Let Ψ_T be a trial state defining a nodal surface, and let $E_{\mathrm{FN}}[\Psi_T]$ be the energy obtained by exact projection constrained to the nodal pockets of Ψ_T. Then

$$E_{\mathrm{FN}}[\Psi_T] \;\geq\; E_0, \tag{51}$$

with equality when Ψ_T reproduces the exact ground-state amplitudes ($\Psi_T = \Psi_0$). The fixed-node energy is therefore a rigorous *upper bound* on the ground-state energy [58]. On a lattice, unlike the continuum, matching only the *nodes* (signs) of Ψ_0 is not sufficient for exactness: both the sign and the relative magnitudes of Ψ_T across every sign-flipping pair of configurations must be correct.

Remark 3.1. The mechanism is variational: the fixed-node procedure yields the exact ground state of an auxiliary Hamiltonian H_{FN} that agrees with H except for an infinite (or, on the lattice, a sign-flip-suppressing) barrier at the imposed nodes. Since H_{FN} is a legitimate constraint restricting the variational manifold, its ground-state energy cannot lie below that of H. Ten Haaf et al. [58] proved this rigorously for lattice fermion Hamiltonians, extending the continuum fixed-node method whose nodal-structure foundation (the tiling theorem) is Ceperley [57].

The determinantal (auxiliary-field) analogue is the *constrained-path Monte Carlo* (CPMC) of Zhang, Carlson, and Gubernatis [59, 60]. In auxiliary-field QMC one propagates a population of Slater determinants $|\phi\rangle$ in imaginary time via a stochastic field $\{s\}$. The sign problem manifests as walkers whose overlap with the trial state, $\langle\Psi_T|\phi\rangle$, changes sign along the random walk. The constrained-path approximation *discards any walker whose overlap turns non-positive*,

$$\langle\Psi_T|\phi\rangle > 0 \quad \text{(constraint retained)}, \qquad \langle\Psi_T|\phi\rangle \leq 0 \quad \text{(walker killed)}, \tag{52}$$

This is analogous to excluding a nodal pocket, but a determinant-overlap boundary in auxiliary-field space is not the same object as a coordinate-space fixed node. In standard CPMC the mixed energy estimator is not generally a rigorous variational upper bound; later analysis corrected that early claim [61]. The method becomes exact for an exact trial state, while a generic approximate Ψ_T produces a systematic bias whose sign is not controlled.

For general Hamiltonians the auxiliary fields (and hence the weights and overlaps) are *complex* rather than merely sign-changing — the Hubbard–Stratonovich decoupling of a general two-body (or *ab initio*) interaction, or a density-channel decoupling of the repulsive interaction away from the half-filled bipartite protection of §3.2.2, produces $e^{i\theta}$ phases. At half filling on a bipartite lattice those density-channel phases pair by the particle–hole argument above; without that pairing one has a genuine phase problem. (The plain repulsive Hubbard model with the discrete *spin*-channel Ising field of Eq. (32) has real ± 1 fields, hence a real *sign* rather than a *phase* problem.) Here the real-axis constraint (52) is insufficient, and one uses the *phaseless* (phase-free) AFQMC of Zhang and Krakauer [62]. The idea is to work with the *importance-sampled* weight referenced to Ψ_T, and to project the complex phase of the incremental overlap ratio onto the real axis via a cosine ("phaseless") projection,

$$w \;\longrightarrow\; w \cdot \max(0,\ \cos\Delta\theta), \qquad \Delta\theta = \arg\frac{\langle\Psi_T|\phi'\rangle}{\langle\Psi_T|\phi\rangle}, \tag{53}$$

so that the walker density is confined near the real, positive region of the overlap. This tames the complex-phase problem at the cost of a trial-state-dependent phaseless bias. Its energy is not generally variational.

Constrained-path and phaseless AFQMC [59, 60, 62] stabilize the random walk by imposing a Ψ_T-defined overlap constraint. Both carry trial-state-dependent bias, and neither has a general variational upper-bound theorem. The rigorous upper bound in Theorem 3.1 belongs to the fixed-node construction.

Honest accounting. Unlike an exact meron cancellation, these constraints trade exponential instability for systematic error with no universal bound. One should vary the trial state, treat back-propagated observables carefully, and benchmark against exact or sign-free limits whenever possible.

3.4 Mitigation by reformulation

For a complex integral, changing the integration manifold can make phase cancellations occur analytically before sampling. The physical partition function is unchanged, but the reference measure—and hence the representation-dependent $\Delta f_{X,N}$ of Eq. (22)—can improve. Such methods generally leave a Jacobian phase, cancellations between sectors, or convergence assumptions, so they are mitigation rather than automatic cures.

3.4.1 Lefschetz thimbles

Consider a bosonic path integral (e.g. after a Hubbard–Stratonovich decoupling) written as an integral over real fields $x \in \mathbb{R}^n$,

$$Z = \int_{\mathbb{R}^n} \mathrm{d}^n x \; e^{-S(x)}, \tag{54}$$

where the action $S(x)$ is complex, so that $\operatorname{Im} S$ oscillates violently along $\mathbb{R}^n$. The central idea, borrowed from the theory of oscillatory integrals and Picard–Lefschetz theory, is to *complexify* the fields, $x \mapsto z \in \mathbb{C}^n$, and to promote S to a holomorphic function $S(z)$ [63, 64]. Because $e^{-S(z)}$ is holomorphic and decays appropriately, Cauchy's theorem lets us deform the real domain $\mathbb{R}^n$ into any homologous n-real-dimensional cycle in $\mathbb{C}^n$ without changing Z.

The distinguished cycles are the *Lefschetz thimbles*. Let z_σ be a critical point (saddle) of the action, $\partial S/\partial z_i\big|_{z_\sigma} = 0$. The thimble $\mathcal{J}_\sigma$ attached to z_σ is the union of all steepest-descent trajectories that flow *into* it (reaching it as $t \to +\infty$) under the *holomorphic gradient flow*

$$\frac{\mathrm{d}z_i}{\mathrm{d}t} = -\overline{\left(\frac{\partial S}{\partial z_i}\right)}, \qquad i = 1, \ldots, n, \tag{55}$$

where the bar denotes complex conjugation and t is the flow (fictitious) time. This flow has two defining properties. Writing $S = \operatorname{Re} S + i \operatorname{Im} S$, one computes along a trajectory

$$\frac{\mathrm{d}S}{\mathrm{d}t} = \sum_i \frac{\partial S}{\partial z_i}\frac{\mathrm{d}z_i}{\mathrm{d}t} = -\sum_i \frac{\partial S}{\partial z_i}\overline{\left(\frac{\partial S}{\partial z_i}\right)} = -\sum_i \left|\frac{\partial S}{\partial z_i}\right|^2 \;\le\; 0, \tag{56}$$

which is real and non-positive. Separating real and imaginary parts,

Along the holomorphic gradient flow (55),

$$\frac{\mathrm{d}(\operatorname{Im} S)}{\mathrm{d}t} = 0, \qquad \frac{\mathrm{d}(\operatorname{Re} S)}{\mathrm{d}t} \le 0. \tag{57}$$

Hence $\operatorname{Im} S$ is *constant* on each thimble (fixed to its saddle value $\operatorname{Im} S(z_\sigma)$), while $\operatorname{Re} S$ decreases monotonically along the flow, attaining its *minimum* at the saddle z_σ (equivalently, $\mathrm{Re}\, S$ grows away from the saddle along $\mathcal{J}_\sigma$). The oscillatory factor $e^{-i\,\mathrm{Im}\, S}$ therefore comes out of the integral over $\mathcal{J}_\sigma$ as a single constant phase.

This is precisely the taming we wanted: on a thimble the integrand is $e^{-i\,\mathrm{Im}\, S(z_\sigma)}\, e^{-\mathrm{Re}\, S}$, and $e^{-\mathrm{Re}\, S}$ is positive and peaked at the saddle. The wildly oscillating measure of (54) has been traded for a smooth, sign-definite one along each descent manifold.

Morse-theory picture. The thimbles are the steepest-descent manifolds attached to the saddles of the Morse function $\mathrm{Re}\, S(z)$ under the downward flow (55): along $\mathcal{J}_\sigma$ the flow descends $\mathrm{Re}\, S$ to its minimum at z_σ. Each critical point contributes one n-dimensional thimble $\mathcal{J}_\sigma$ and one dual, steepest-*ascent* cycle $\mathcal{K}_\sigma$. Picard–Lefschetz theory expresses the original contour as an integer combination of thimbles,

$$Z = \sum_\sigma m_\sigma\, e^{-i\,\mathrm{Im}\, S(z_\sigma)} \int_{\mathcal{J}_\sigma} \mathrm{d}^n z \;\, e^{-\mathrm{Re}\, S(z)}, \qquad m_\sigma = \langle \mathbb{R}^n, \mathcal{K}_\sigma \rangle \in \mathbb{Z}, \tag{58}$$

where the intersection numbers m_σ count how the real cycle threads the ascending manifolds [63].

The residual sign. Two effects prevent this from being a complete cure. First, parametrizing a thimble by real coordinates introduces a *Jacobian* $\det J$ (from the tangent-space frame transported by the flow); $\det J$ is in general complex, so a *residual phase* $e^{i\arg\det J}$ remains inside each thimble integral. Second, when several thimbles contribute ($m_\sigma \neq 0$ for more than one σ), their constant phases $e^{-i\,\mathrm{Im}\, S(z_\sigma)}$ differ and *cancel against one another*. A dominant thimble often helps empirically, but there is no general bound guaranteeing that its Jacobian phase is mild [5].

Practical manifolds. Flowing exactly onto a thimble is numerically stiff. The first Monte Carlo simulations performed *on* a thimble demonstrated that the geometric proposal can be turned into a working method [65]. A productive compromise is to flow only *partway* ($t < \infty$), landing on a manifold that is not a thimble but already has strongly suppressed phase fluctuations [64]; the residual phase is then treated by reweighting. This motivates *sign-optimized manifolds*, in which the deformation is parametrized and its parameters tuned to maximize the average sign [66], together with the closely related path-optimization method [67]; a tempered-flow variant uses the flow time itself as a tempering parameter to reach the doped Hubbard model [68].

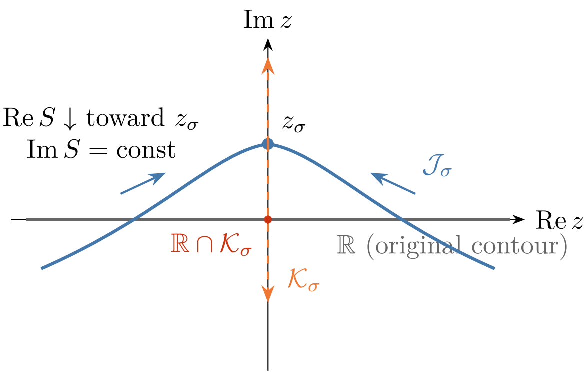

Figure 5: One-thimble schematic for $n = 1$. Downward flow on $\mathcal{J}_\sigma$ approaches the saddle, keeping $\operatorname{Im} S$ constant while decreasing $\operatorname{Re} S$. The dual cycle $\mathcal{K}_\sigma$ fixes the intersection number m_σ. In general the real contour is homologous to a *sum* $\sum_\sigma m_\sigma \mathcal{J}_\sigma$, and a complex Jacobian can still supply a residual phase on each thimble.

3.4.2 Complex Langevin

A conceptually distinct route to the complexified configuration space is *complex Langevin* (CL), based on stochastic quantization; its modern revival as a sign-problem tool dates to Aarts [69], who showed on the relativistic Bose gas at finite chemical potential — a theory with a genuine sign/Silver-Blaze problem — that the complexified dynamics can reproduce the correct thermodynamics where reweighting fails. For a real theory, expectation values can be computed as long-time averages of a Langevin process

$$\frac{\mathrm{d}x_i}{\mathrm{d}\tau} = -\frac{\partial S}{\partial x_i} + \eta_i(\tau), \qquad \langle \eta_i(\tau)\eta_j(\tau')\rangle = 2\,\delta_{ij}\,\delta(\tau-\tau'), \tag{59}$$

whose stationary distribution is $\propto e^{-S(x)}$. When S is complex the drift $-\partial S/\partial x_i$ is complex, so the trajectory leaves the real axis. The remedy is to let the fields complexify, $x_i \to z_i = x_i + iy_i$, and evolve

$$\frac{\mathrm{d}z_i}{\mathrm{d}\tau} = -\frac{\partial S}{\partial z_i} + \eta_i(\tau), \tag{60}$$

with real noise η_i (a complex-noise variant also exists). One then measures holomorphic observables $O(z)$ along the process. Unlike thimbles, CL requires *no explicit knowledge of saddles or contours*: the stochastic dynamics finds its own equilibrium region in $\mathbb{C}^n$, which makes it attractive for high-dimensional lattice actions [6].

Remark 3.2 (Honest caveat). CL is powerful but *correct convergence is not guaranteed*. The formal proof that the CL average reproduces the correct expectation values relies on assumptions that can fail in practice. Two well-documented failure modes are: (i) *boundary terms* — if the probability distribution of the complexified fields does not decay fast enough at large $|y_i|$, integration by parts in the correctness proof picks up nonzero surface contributions and CL converges to the *wrong* answer; and (ii) *poles / singular drift* — when S contains logarithms of a fermion determinant, $-\partial S/\partial z$ has poles wherever $\det = 0$, and trajectories orbiting these singularities spoil ergodicity and correctness [6]. Diagnostics (monitoring the decay of the drift-magnitude distribution) and remedies (dynamical stabilization, gauge cooling, kernel/regulator modifications) exist, but one must always *verify* convergence against a known limit rather than assume it.

3.4.3 The current frontier: learned contour deformations

Learned manifolds are trainable contour deformations, not complex Langevin. Write a holomorphic deformation as $z = F_\vartheta(x)$, with $x \in \mathbb{R}^n$. Sampling in x uses the effective action

$$S_{\text{eff}}(x;\vartheta) = S(F_\vartheta(x)) - \log\det J_\vartheta(x), \qquad J_\vartheta = \frac{\partial F_\vartheta}{\partial x}. \tag{61}$$

The parameters can be trained to increase the average phase (or minimize a phase-quenched loss), using normalizing flows or more general neural maps [5, 70]. Training does not guarantee a global optimum: multimodal thimble structure, expensive Jacobians, zeros of the fermion determinant, and residual Jacobian phases can all remain. Learned contours are therefore adaptive mitigation with diagnostics, not an automatic sign cure. For a broader review of contour and Langevin methods see Aarts and Sexty [7].

3.5 Fundamental limit I: generic efficient curing is NP-hard

The Troyer–Wiese result concerns a *polynomial-time simulation*, not the mere existence of non-negative weights [23]. This distinction is essential. A basis change can move a sign problem into an NP-hard positive ensemble with exponentially slow mixing; that has not solved the computational problem.

Following the original formulation, consider a signed QMC representation whose absolute-value (bosonic) reference problem can already be simulated in polynomial time. A solution of its sign problem is an unbiased estimator of the required thermal observable whose total cost for fixed relative accuracy is polynomial in system size and inverse temperature. The total cost includes autocorrelation time and the effort needed to construct the representation.

Theorem 3.2 (Troyer–Wiese, informal)**.** The task of providing such a polynomial-time solution for arbitrary signed QMC instances is NP-hard. A generic deterministic solution would imply P = NP; a randomized solution would place the corresponding NP-complete decision problem in bounded-error randomized polynomial time.

3.5.1 The signed spin-glass construction

Start with the three-dimensional Ising spin-glass threshold problem

$$H_{\mathrm{SG}}(s) = -\sum_{\langle ij\rangle} J_{ij} s_i s_j, \qquad s_i = \pm 1, \qquad J_{ij} \in \{+J, -J\}. \tag{62}$$

Deciding whether some configuration has energy at most E_{th} is NP-complete [71]. Encode the same spectrum in the quantum Hamiltonian

$$H_X = -\sum_{\langle ij\rangle} J_{ij} \sigma_i^x \sigma_j^x. \tag{63}$$

The simultaneous σ^x eigenstates reproduce Eq. (62), but QMC in the σ^z product basis sees off-diagonal matrix elements $-J_{ij}$ of both signs. Equivalently, the SSE vertices of $-H_X$ carry the random coefficients J_{ij}, so frustrated operator strings have signed weights.

Taking absolute values replaces every J_{ij} by $|J_{ij}| = J$. The reference Hamiltonian is therefore the unfrustrated ferromagnet

$$H_{\mathrm{ref}} = -J \sum_{\langle ij\rangle} \sigma_i^x \sigma_j^x, \tag{64}$$

for which efficient cluster algorithms are available. This is why the constructed instance isolates the sign cancellation rather than hiding an already-hard reference ensemble.

The spectrum of Eq. (62) has a fixed energy spacing of order J and at most 2^N states. Choosing $\beta = O[(N + \log(1/\epsilon))/J]$ suppresses the total excited-state contribution sufficiently that an estimate of the thermal energy to fixed accuracy resolves the ground-state threshold. A hypothetical generic polynomial-time cure for the signed representation of H_X would therefore

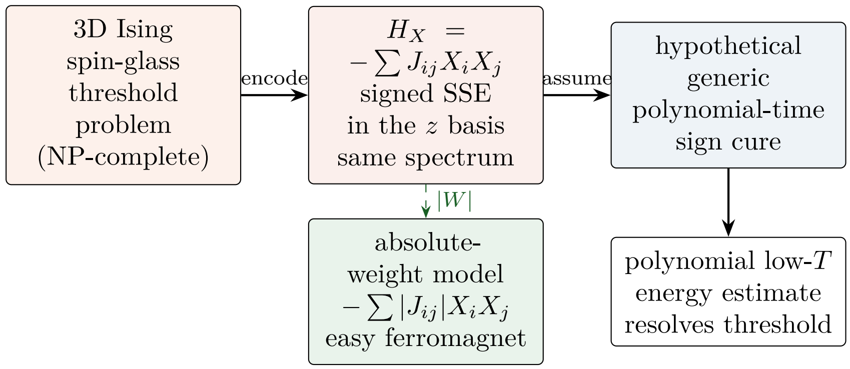

Figure 6: The Troyer–Wiese reduction. Random spin-glass couplings become signed off-diagonal vertices in the z-basis QMC representation of H_X, while the absolute-weight reference is an easy ferromagnet. A generic cure with polynomial total sampling cost would estimate the low-temperature energy and solve the NP-complete threshold problem.

decide the NP-complete spin-glass problem in polynomial time. The reduction is summarized in Fig. 6.

A local Hadamard rotation maps H_X to the diagonal classical Hamiltonian $-\sum J_{ij}\sigma_i^z\sigma_j^z$. This removes the matrix-element signs but leaves the NP-hard glassy landscape. The example is a useful stress test for terminology: *sign-free* and *efficiently simulable* are not synonyms.

Remark 3.3 (Worst-case scope). The theorem does not say that every sign problem, or any particular material model, is NP-hard. It excludes a uniform polynomial-time method for all signed instances. Symmetry-based determinant pairing, meron cancellations, and other model-specific constructions evade the reduction by restricting the input class.

3.6 Fundamental limit II: finding a stoquastic basis can be hard

Representation dependence suggests an automated search: given a local Hamiltonian, choose on-site unitaries $U = \bigotimes_i u_i$ so that $UHU^\dagger$ is stoquastic. The complexity of this question depends on the allowed transformations and on whether stoquasticity is required term by term or only after summing local terms. It is therefore not one undifferentiated decision problem.

Theorem 3.3 (Restricted curing is NP-complete)**.** For constructed families of local qubit Hamiltonians, deciding whether a cure exists is NP-complete when the allowed transformations are single-qubit Clifford operations or single-qubit real orthogonal matrices [72].

This result encodes a global constraint-satisfaction problem: a rotation that fixes one interaction can spoil another interaction sharing the same qubit. It does not establish NP membership for every formulation involving arbitrary continuous unitaries.

For exactly two-local qubit Hamiltonians, a sharper structural distinction is known [73]:

- If the Hamiltonian contains only two-qubit terms and no one-qubit fields, curability by single-qubit rotations can be decided constructively using polynomially many arithmetic operations over $\mathbb{R}$.
- Once one-qubit terms are allowed, the task can be NP-hard. Klassen *et al.* construct a restricted family for which curing is equivalent to 3-SAT and hence NP-complete.

The fields pin local frames, allowing shared two-local constraints (with auxiliary qubits in the reduction) to encode Boolean clauses. There is no literal three-body $X_i X_j X_k$ clause in this two-local theorem. The source-faithful conclusion is the complexity dichotomy above, not a generic claim that every Hamiltonian with fields is hard.

Even optimizing a quantitative measure of positive off-diagonal weight can be NP-complete for specified local basis families, by a reduction related to MaxCut [74]. Thus both perfect curing and systematic easing have hard worst cases. These statements motivate the constructive certificates in Sec. 3.2: they verify membership in a positive subclass without searching the full space of local frames. As before, worst-case hardness neither rules out a cure for a particular physical model nor proves that a model without a known cure is intrinsically sign-problematic.

3.7 Fundamental limit III: intrinsic obstructions

Worst-case hardness asks whether one algorithm works on every input. An *intrinsic sign problem* asks a different, phase-level question: can any representative of a phase be made positive within a specified class of local Hamiltonians and Monte Carlo reformulations? The qualifier is indispensable; no cited theorem rules out literally every imaginable resummation.

3.7.1 Three levels of obstruction

It is useful to separate three statements.

(S1) A chosen basis or auxiliary-field decomposition has negative or complex weights.

(S2) A fixed Hamiltonian cannot be made stoquastic by a specified family of transformations, such as on-site rotations.

(S3) A whole phase has no positive representative within a specified local algorithmic class, certified by a phase invariant.

In standard usage, *non-stoquastic* describes a Hamiltonian in a specified basis; it does not already include the quantifier in (S2). Moreover, (S1) does not imply (S2), and failure to find a cure for (S2) does not prove (S3). Figure 7 displays the correct logical workflow.

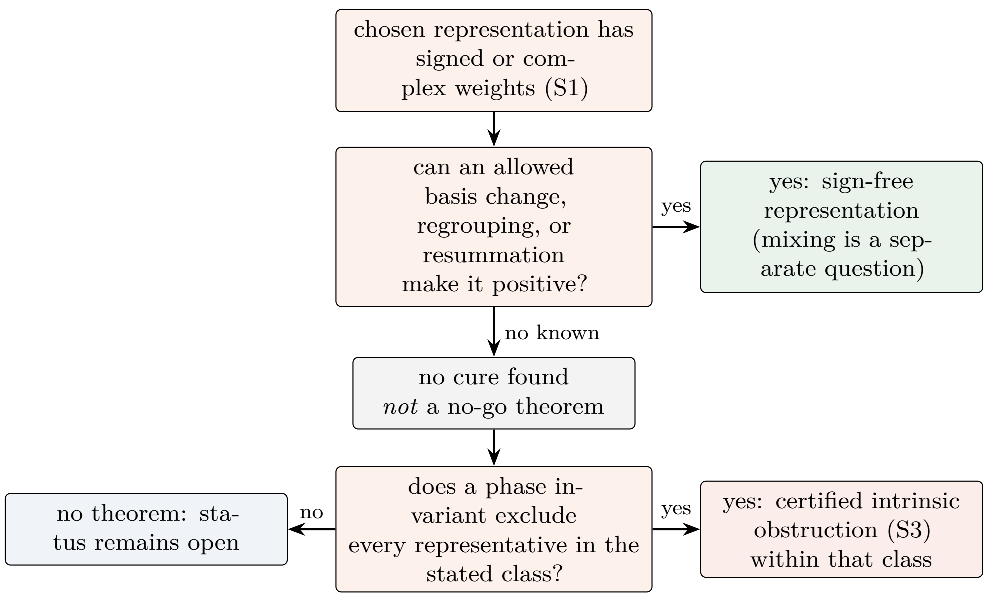

Figure 7: A sign in one formulation is only the start of the analysis. Non-stoquastic Hamiltonians can possess positive regrouped expansions; twisted quantum doubles are an example. Intrinsic status requires a theorem, with an explicitly stated class of allowed representatives and transformations.

3.7.2 Restricted wavefunction and response obstructions

Perron–Frobenius implies that a stoquastic Hamiltonian has a non-negative ground state in each connected block. The converse strategy is therefore to show that a phase cannot support such a ground state under the allowed local changes. Hastings proved a restricted result for commuting Hamiltonians in the double-semion phase under the technical TQO-2 assumption [75]. It excludes a non-negative commuting-projector representative in that setting; it is not a theorem against every positive regrouped SSE representation.

Quantized gravitational response supplies a more physical obstruction. For a two-dimensional gapped chiral phase,

$$\kappa_{xy} = c_- \frac{\pi^2 k_B^2}{3h} T = c_- \frac{\pi k_B^2}{6\hbar} T, \tag{65}$$

where c_- is the chiral central charge. Ringel and Kovrizhin showed that quantized gravitational responses obstruct a class of local sign-free QMC representations for bosonic systems [76]. The result has locality and representation assumptions; it should not be paraphrased as an unconditional theorem that every $c_- \neq 0$ phase defeats every possible algorithm.

A sharper chiral criterion combines c_- with the anyon topological spins $\{\theta_a\}$. Within the local-stoquastic and local sign-free-DQMC framework of Golan, Smith, and Ringel,

$$\boxed{e^{2\pi i c_-/24} \notin \{\theta_a\} \quad \Longrightarrow \quad \text{a local positivity obstruction.}} \tag{66}$$

The bosonic theorem excludes locally stoquastic representatives; its fermionic counterpart excludes the local sign-free determinant classes considered in that work [77]. Nonzero c_- alone is a useful warning sign, not the precise criterion in Eq. (66).

For broad classes of bosonic $2+1$-dimensional topological field theories, the anyon data provide another sufficient diagnostic: if the exchange statistics cannot be arranged into the required complete sets of roots of unity, a local stoquastic representative is obstructed [78]. This criterion is sharpest in the gapped Abelian setting treated by the theorem; stronger simultaneous-positivity tests and non-Abelian extensions have additional assumptions. Recent higher-Gauss-sum diagnostics continue to refine the boundary of the obstructed class [79].

3.7.3 Why non-stoquasticity is not intrinsic

The 2026 construction of Shackleton is a decisive counterexample to an overly broad identification of signs with intrinsic obstruction. The double-semion model and, more generally, twisted quantum doubles for finite groups admit local Hamiltonians whose standard matrix representation is non-stoquastic but whose SSE weights become non-negative after the cocycle phases are grouped around closed histories [80]. This does not contradict the restricted commuting-projector/non-negative-wavefunction theorem above; it changes the allowed positive representation.

The lesson is methodological. Statements about intrinsic signs must specify at least

- whether the objects being classified are Hamiltonian matrix elements, ground-state amplitudes, worldline weights, or determinants;
- which changes of basis, finite-depth circuits, auxiliary fields, and regroupings are allowed; and
- whether the conclusion concerns one Hamiltonian or every representative along a gapped path in a phase.

3.7.4 Open cases and the role of the sign

Known intrinsic examples form a restricted and incompletely classified set. The sign problems that dominate practice are mostly not known to be intrinsic: the doped two-dimensional Hubbard model, fermions at baryon chemical potential, and real-time dynamics remain targets for new representations. Finite density itself is not synonymous with a phase problem—isospin chemical potential and two-color gauge theories, for example, retain special positivity—and doped Hubbard DQMC usually has real signed rather than complex weights.

Within a fixed formulation, the average sign can still be informative. Its finite-size scaling may sharpen near a phase transition, and rigorous bounds can relate it to the spectrum of a chosen positive reference problem [21, 22, 81]. Because $\langle \mathrm{sgn} \rangle$ changes under a reformulation, it is a representation-dependent diagnostic, not a conventional basis-independent observable.

Quantum processors may bypass positive classical weights for selected real-time evolutions and quantum-simulation tasks. They do not automatically prepare low-temperature states, compute partition functions, or solve the NP-hard instances used in Sec. 3.5. Quantum hardware is therefore a complementary route for particular observables, not a universal resolution of the equilibrium sign problem.

4 Conclusion

The sign problem is best understood as a failure of a chosen classical sampling representation. In SSE it appears in closed products of Hamiltonian matrix elements; in worldline language it records permutation parity; in auxiliary-field QMC it is compressed into a determinant. Reweighting is exact, and a complex representation of a real partition function can be conjugate-paired into the real weight $\operatorname{Re} W$. The resulting $|\operatorname{Re} W|$ reference has an average sign at least as large as the conventional average phase obtained from $|W|$. When the relevant Δf_N approaches a positive constant, the sign-induced cost grows as $\exp(2\beta N \Delta f_N)$; when $\Delta f_N \sim (\log N)/(\beta N)$, it is only algebraic.

Exact progress comes from structure: Marshall rotations, determinant and Kramers pairing, fermion-bag and meron resummations, Majorana positivity, and split-orthogonal orientation theorems. Their common achievement is positivity, not an automatic proof of rapid mixing. Approximate methods and contour deformations can reach farther, but their residual phases, convergence conditions, or trial-state biases must be reported explicitly.

The fundamental results have equally distinct scopes. Troyer–Wiese rules out a generic efficient cure in the worst case; basis-search results make some automated local cures hard; topological response and anyon data exclude positive representatives only within stated local algorithmic classes. Keeping these quantifiers visible turns a collection of apparent contradictions into a coherent map of what is known, what is model-specific, and what remains open.

Acknowledgments

The author acknowledges the use of OpenAI Codex with the GPT-5.6 model for assistance in reorganizing and editing these notes, checking internal consistency, and improving the presentation of figures.

References

[1] Gaopei Pan and Zi Yang Meng. Sign problem in quantum monte carlo simulation. *arXiv preprint arXiv:2204.08777*, 2022.

[2] Xiao-Yan Xu. Quantum monte carlo study of strongly correlated electrons. *Acta Physica Sinica*, 71(12):127101, 2022. doi: 10.7498/aps.71.20220079. URL https://doi.org/10.7498/aps.71.20220079.

[3] F. F. Assaad and H. G. Evertz. World-line and determinantal quantum monte carlo methods for spins, phonons and electrons. In H. Fehske, R. Schneider, and A. Weiße, editors, *Computational Many-Particle Physics*, volume 739 of *Lecture Notes in Physics*, pages 277–356. Springer, Berlin, Heidelberg, 2008. doi: 10.1007/978-3-540-74686-7_10. URL https://doi.org/10.1007/978-3-540-74686-7_10.

[4] Zi-Xiang Li and Hong Yao. Sign-problem-free fermionic quantum monte carlo: Developments and applications. *Annual Review of Condensed Matter Physics*, 10(1):337–356, 2019. ISSN 1947-5454. doi: 10.1146/annurev-conmatphys-033117-054307. URL https://doi.org/10.1146/annurev-conmatphys-033117-054307.

[5] Andrei Alexandru, Gökçe Başar, Paulo F. Bedaque, and Neill C. Warrington. Complex paths around the sign problem. *Reviews of Modern Physics*, 94(1):015006, 2022. doi: 10.1103/RevModPhys.94.015006. URL https://link.aps.org/doi/10.1103/RevModPhys.94.015006.

[6] C. E. Berger, L. Rammelmueller, A. C. Loheac, F. Ehmann, J. Braun, and J. E. Drut. Complex langevin and other approaches to the sign problem in quantum many-body physics. *Physics Reports-Review Section of Physics Letters*, 892:1–54, 2021. ISSN 0370-1573. doi: 10.1016/j.physrep.2020.09.002. URL https://doi.org/10.1016/j.physrep.2020.09.002.

[7] Gert Aarts and Denes Sexty. Lattice field theories with a sign problem. *arXiv preprint arXiv:2604.24290*, 2026. URL https://arxiv.org/abs/2604.24290.

[8] Anders W. Sandvik. Stochastic series expansion method with operator-loop update. *Physical Review B*, 59(22):R14157–R14160, 1999. doi: 10.1103/PhysRevB.59.R14157. URL https://link.aps.org/doi/10.1103/PhysRevB.59.R14157.

[9] Anders W. Sandvik. Computational studies of quantum spin systems. *AIP Conference Proceedings*, 1297(1):135–338, 2010. doi: 10.1063/1.3518900. URL https://doi.org/10.1063/1.3518900.

[10] W. Marshall. Antiferromagnetism. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 232 (1188):48–68, 1955. doi: 10.1098/rspa.1955.0200. URL https://royalsocietypublishing.org/doi/10.1098/rspa.1955.0200.

[11] H. G. Evertz. The loop algorithm. *Advances in Physics*, 52(1):1–66, 2003. ISSN 0001-8732. doi: 10.1080/0001873021000049195. URL http://dx.doi.org/10.1080/0001873021000049195.

[12] Ribhu K. Kaul. Marshall-positive su(n) quantum spin systems and classical loop models: A practical strategy to design sign-problem-free spin hamiltonians. *Physical Review B*, 91(5), 2015. ISSN 1098-0121. doi: 10.1103/PhysRevB.91.054413. URL https://doi.org/10.1103/PhysRevB.91.054413.

[13] Ribhu K. Kaul, Roger G. Melko, and Anders W. Sandvik. Bridging lattice-scale physics and continuum field theory with quantum monte carlo simulations. *Annual Review of Condensed Matter Physics*, 4(1):179–215, 2013. ISSN 1947-5454. doi: 10.1146/annurev-conmatphys-030212-184215. URL https://doi.org/10.1146/annurev-conmatphys-030212-184215.

[14] Fabien Alet, Kedar Damle, and Sumiran Pujari. Sign-problem-free monte carlo simulation of certain frustrated quantum magnets. *Physical Review Letters*, 117(19), 2016. ISSN 0031-9007. doi: 10.1103/PhysRevLett.117.197203. URL https://doi.org/10.1103/PhysRevLett.117.197203.

[15] Stefan Wessel, Ido Niesen, Jonas Stapmanns, B. Normand, Frederic Mila, Philippe Corboz, and Andreas Honecker. Thermodynamic properties of the shastry-sutherland model from quantum monte carlo simulations. *Physical Review B*, 98(17), 2018. ISSN 2469-9950. doi: 10.1103/PhysRevB.98.174432. URL https://doi.org/10.1103/PhysRevB.98.174432.

[16] P. Henelius and A. W. Sandvik. Sign problem in monte carlo simulations of frustrated quantum spin systems. *Physical Review B*, 62(2):1102–1113, 2000. ISSN 1098-0121. doi: 10.1103/PhysRevB.62.1102. URL https://doi.org/10.1103/PhysRevB.62.1102.

[17] A. Yu. Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1):2–30, 2003. doi: 10.1016/S0003-4916(02)00018-0. URL https://doi.org/10.1016/S0003-4916(02)00018-0.

[18] Chaithanya Rayudu and Jun Takahashi. Fast mixing of operator-loop path-integral quantum monte carlo for stoquastic xy hamiltonians. *arXiv preprint arXiv:2509.21683*, 2025. URL https://arxiv.org/abs/2509.21683.

[19] Attila Pasztor, Szabolcs Borsanyi, Zoltan Fodor, Kornel Kapas, Sandor D. Katz, Matteo Giordano, Daniel Nogradi, and Chik Him Wong. New approach to lattice qcd at finite density: reweighting without an overlap problem. *arXiv preprint arXiv:2112.02134*, 2021. URL https://arxiv.org/abs/2112.02134.

[20] Yunqing Ouyang and Xiao Yan Xu. Projection of infinite-U hubbard model and algebraic sign structure. *Physical Review B*, 104(24):L241104, 2021. doi: 10.1103/PhysRevB.104.L241104. URL https://link.aps.org/doi/10.1103/PhysRevB.104.L241104.

[21] Xu Zhang, Gaopei Pan, Xiao Yan Xu, and Zi Yang Meng. Fermion sign bounds theory in quantum monte carlo simulation. *Physical Review B*, 106(3):035121, 2022. doi: 10.1103/PhysRevB.106.035121. URL https://link.aps.org/doi/10.1103/PhysRevB.106.035121.

[22] R. Mondaini, S. Tarat, and R. T. Scalettar. Quantum critical points and the sign problem. *Science*, 375(6579):418–424, 2022. doi: 10.1126/science.abg9299. URL https://doi.org/10.1126/science.abg9299.

[23] Matthias Troyer and Uwe-Jens Wiese. Computational complexity and fundamental limitations to fermionic quantum monte carlo simulations. *Physical Review Letters*, 94(17), 2005. ISSN 0031-9007 1079-7114. doi: 10.1103/PhysRevLett.94.170201.

[24] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar. Sign problem in the numerical simulation of many-electron systems. *Physical Review B*, 41(13):9301–9307, 1990. doi: 10.1103/PhysRevB.41.9301. URL https://link.aps.org/doi/10.1103/PhysRevB.41.9301.

[25] R. Blankenbecler, D. J. Scalapino, and R. L. Sugar. Monte carlo calculations of coupled boson-fermion systems. i. *Physical Review D*, 24(8):2278–2286, 1981. doi: 10.1103/PhysRevD.24.2278. URL https://link.aps.org/doi/10.1103/PhysRevD.24.2278.

[26] G. Sugiyama and S. E. Koonin. Auxiliary field monte-carlo for quantum many-body ground states. *Annals of Physics*, 168(1):1–26, 1986. ISSN

0003-4916. doi: 10.1016/0003-4916(86)90107-7. URL https://doi.org/10.1016/0003-4916(86)90107-7.

[27] S. Sorella, S. Baroni, R. Car, and M. Parrinello. A novel technique for the simulation of interacting fermion systems. *EPL (Europhysics Letters)*, 8(7):663, 1989. URL http://stacks.iop.org/0295-5075/8/i=7/a=014.

[28] J. E. Hirsch. Discrete hubbard-stratonovich transformation for fermion lattice models. *Physical Review B*, 28(7):4059–4061, 1983. doi: 10.1103/PhysRevB.28.4059. URL https://link.aps.org/doi/10.1103/PhysRevB.28.4059.

[29] Liang Chen and A. M. S. Tremblay. Determinant monte carlo for the hubbard model with arbitrarily gauged auxiliary fields. *International Journal of Modern Physics B*, 06(05n06):547–560, 1992. ISSN 0217-9792. doi: 10.1142/S0217979292000323. URL https://doi.org/10.1142/S0217979292000323.

[30] J. E. Hirsch, D. J. Scalapino, R. L. Sugar, and R. Blankenbecler. Efficient monte carlo procedure for systems with fermions. *Physical Review Letters*, 47(22):1628–1631, 1981. doi: 10.1103/PhysRevLett.47.1628. URL https://link.aps.org/doi/10.1103/PhysRevLett.47.1628.

[31] S. R. White, D. J. Scalapino, R. L. Sugar, E. Y. Loh, J. E. Gubernatis, and R. T. Scalettar. Numerical study of the two-dimensional hubbard model. *Physical Review B*, 40(1):506–516, 1989. doi: 10.1103/PhysRevB.40.506. URL https://link.aps.org/doi/10.1103/PhysRevB.40.506.

[32] V. I. Iglovikov, E. Khatami, and R. T. Scalettar. Geometry dependence of the sign problem in quantum monte carlo simulations. *Physical Review B*, 92(4), 2015. ISSN 2469-9950. doi: 10.1103/PhysRevB.92.045110. URL https://doi.org/10.1103/PhysRevB.92.045110.

[33] Elliott H. Lieb. Two theorems on the hubbard model. *Physical Review Letters*, 62(10):1201–1204, 1989. doi: 10.1103/PhysRevLett.62.1201. URL https://link.aps.org/doi/10.1103/PhysRevLett.62.1201. Erratum: Phys. Rev. Lett. 62, 1927 (1989).

[34] G. S. Tian. Lieb's spin-reflection-positivity method and its applications to strongly correlated electron systems. *Journal of Statistical Physics*, 116(1-4):629–680, 2004. ISSN 0022-4715. doi: 10.1023/b:

Joss.0000037214.70064.78. URL https://doi.org/10.1023/b:Joss.0000037214.70064.78.

[35] C. J. Wu and S. C. Zhang. Sufficient condition for absence of the sign problem in the fermionic quantum monte carlo algorithm. *Physical Review B*, 71(15), 2005. ISSN 2469-9950. doi: ARTN15511510.1103/PhysRevB.71.155115. URL https://doi.org/10.1103/PhysRevB.71.155115.

[36] Erez Berg, Max A. Metlitski, and Subir Sachdev. Sign-problem-free quantum monte carlo of the onset of antiferromagnetism in metals. *Science*, 338(6114):1606–1609, 2012. ISSN 0036-8075. doi: 10.1126/science.1227769. URL https://doi.org/10.1126/science.1227769.

[37] Yoni Schattner, Samuel Lederer, Steven A. Kivelson, and Erez Berg. Ising nematic quantum critical point in a metal: A monte carlo study. *Physical Review X*, 6(3):031028, 2016. doi: 10.1103/PhysRevX.6.031028. URL https://link.aps.org/doi/10.1103/PhysRevX.6.031028.

[38] Max H. Gerlach, Yoni Schattner, Erez Berg, and Simon Trebst. Quantum critical properties of a metallic spin-density-wave transition. *Physical Review B*, 95(3):035124, 2017. doi: 10.1103/PhysRevB.95.035124. URL https://link.aps.org/doi/10.1103/PhysRevB.95.035124.

[39] Shailesh Chandrasekharan. Fermion bag approach to lattice field theories. *Physical Review D*, 82(2), 2010. ISSN 1550-7998. doi: 10.1103/PhysRevD.82.025007. URL https://doi.org/10.1103/PhysRevD.82.025007.

[40] Emilie Fulton Huffman and Shailesh Chandrasekharan. Solution to sign problems in half-filled spin-polarized electronic systems. *Physical Review B*, 89(11):111101, 2014. doi: 10.1103/PhysRevB.89.111101. URL https://link.aps.org/doi/10.1103/PhysRevB.89.111101.

[41] Lei Wang, Philippe Corboz, and Matthias Troyer. Fermionic quantum critical point of spinless fermions on a honeycomb lattice. *New Journal of Physics*, 16(10):103008, 2014. ISSN 1367-2630. doi: 10.1088/1367-2630/16/10/103008.

[42] Shailesh Chandrasekharan. Fermion bag approach to fermion sign problems. *The European Physical Journal A*, 49(7):90, 2013. ISSN 1434-601X. doi: 10.1140/epja/i2013-13090-y. URL https://doi.org/10.1140/epja/i2013-13090-y.

[43] Emilie Huffman and Shailesh Chandrasekharan. Solution to new sign problems with hamiltonian lattice fermions. *arXiv preprint arXiv:1411.7147*, 2014.

[44] Connor T. Hann, Emilie Huffman, and Shailesh Chandrasekharan. Solution to the sign problem in a frustrated quantum impurity model. *Annals of Physics*, 376:63–75, 2017. ISSN 0003-4916. doi: 10.1016/j.aop.2016.11.006. URL https://doi.org/10.1016/j.aop.2016.11.006.

[45] Zi-Xiang Li, Yi-Fan Jiang, and Hong Yao. Solving the fermion sign problem in quantum monte carlo simulations by majorana representation. *Physical Review B*, 91(24):241117, 2015. URL http://link.aps.org/doi/10.1103/PhysRevB.91.241117.

[46] Zi-Xiang Li, Yi-Fan Jiang, and Hong Yao. Majorana-time-reversal symmetries: A fundamental principle for sign-problem-free quantum monte carlo simulations. *Physical Review Letters*, 117(26):267002, 2016. URL http://link.aps.org/doi/10.1103/PhysRevLett.117.267002.

[47] Z. C Wei, Congjun Wu, Yi Li, Shiwei Zhang, and T. Xiang. Majorana positivity and the fermion sign problem of quantum monte carlo simulations. *Physical Review Letters*, 116(25):250601, 2016. URL http://link.aps.org/doi/10.1103/PhysRevLett.116.250601.

[48] Tomoya Hayata and Arata Yamamoto. Classification of sign-problem-free relativistic fermions on the basis of the majorana positivity. *Progress of Theoretical and Experimental Physics*, 2017(4), 2017. ISSN 2050-3911. doi: 10.1093/ptep/ptx040. URL https://doi.org/10.1093/ptep/ptx040.

[49] Zhong-Chao Wei. Semigroup approach to the sign problem in quantum monte carlo simultions. *arXiv preprint arXiv:1712.09412*, 2017.

[50] Lei Wang, Ye-Hua Liu, Mauro Iazzi, Matthias Troyer, and Gergely Harcos. Split orthogonal group: A guiding principle for sign-problem-free fermionic simulations. *Physical Review Letters*, 115(25):250601, 2015. URL http://link.aps.org/doi/10.1103/PhysRevLett.115.250601.

[51] Xiao Yan Xu, Yang Qi, Long Zhang, Fakher F. Assaad, Cenke Xu, and Zi Yang Meng. Monte carlo study of lattice compact quantum electrodynamics with fermionic matter: The parent state of quantum phases.

Physical Review X, 9(2):021022, 2019. doi: 10.1103/PhysRevX.9.021022. URL https://link.aps.org/doi/10.1103/PhysRevX.9.021022.

[52] Xu Zhang and Nick Bultinck. Constraints from antiunitary symmetries on phase diagrams of sign problem free models. *Physical Review B*, 112 (12):125110, 2025. doi: 10.1103/pxsx-33qy. URL https://link.aps.org/doi/10.1103/pxsx-33qy.

[53] Ori Grossman and Erez Berg. Unavoidable fermi liquid instabilities in sign problem-free models. *arXiv preprint arXiv:2302.11262*, 2023.

[54] S. Chandrasekharan and U. J. Wiese. Meron-cluster solution of fermion sign problems. *Physical Review Letters*, 83(16):3116–3119, 1999. ISSN 0031-9007. doi: 10.1103/PhysRevLett.83.3116. URL https://doi.org/10.1103/PhysRevLett.83.3116.

[55] S. Bergkvist, P. Henelius, and A. Rosengren. Reduction of the sign problem using the meron-cluster approach. *Physical Review E*, 68 (1), 2003. ISSN 1539-3755. doi: 10.1103/PhysRevE.68.016122. URL https://doi.org/10.1103/PhysRevE.68.016122.

[56] S. Chandrasekharan, J. Cox, K. Holland, and U. J. Wiese. Meron-cluster simulation of a chiral phase transition with staggered fermions. *Nuclear Physics B*, 576(1-3):481–500, 2000. ISSN 0550-3213. doi: 10.1016/s0550-3213(00)00087-0. URL https://doi.org/10.1016/s0550-3213(00)00087-0.

[57] D. M. Ceperley. Fermion nodes. *Journal of Statistical Physics*, 63(5-6): 1237–1267, 1991. ISSN 0022-4715. doi: 10.1007/BF01030009. URL https://doi.org/10.1007/BF01030009.

[58] D. F. B. ten Haaf, H. J. M. van Bemmel, J. M. J. van Leeuwen, W. van Saarloos, and D. M. Ceperley. Proof for an upper bound in fixed-node monte carlo for lattice fermions. *Physical Review B*, 51(19):13039–13045, 1995. doi: 10.1103/PhysRevB.51.13039. URL https://link.aps.org/doi/10.1103/PhysRevB.51.13039.

[59] S. W. Zhang, J. Carlson, and J. E. Gubernatis. Constrained path monte carlo method for fermion ground states. *Physical Review B*, 55(12): 7464–7477, 1997. ISSN 0163-1829. doi: 10.1103/PhysRevB.55.7464. URL https://doi.org/10.1103/PhysRevB.55.7464.

[60] Shiwei Zhang, J. Carlson, and J. E. Gubernatis. Constrained path quantum monte carlo method for fermion ground states. *Physical Review Letters*, 74(18):3652–3655, 1995. doi: 10.1103/PhysRevLett.74.3652. URL https://link.aps.org/doi/10.1103/PhysRevLett.74.3652.

[61] J. Carlson, J. E. Gubernatis, G. Ortiz, and Shiwei Zhang. Issues and observations on applications of the constrained-path monte carlo method to many-fermion systems. *Physical Review B*, 59(19):12788–12798, 1999. doi: 10.1103/PhysRevB.59.12788. URL https://doi.org/10.1103/PhysRevB.59.12788.

[62] S. W. Zhang and H. Krakauer. Quantum monte carlo method using phase-free random walks with slater determinants. *Physical Review Letters*, 90(13), 2003. ISSN 0031-9007. doi: 10.1103/PhysRevLett.90.136401. URL https://doi.org/10.1103/PhysRevLett.90.136401.

[63] Marco Cristoforetti, Francesco Di Renzo, Luigi Scorzato, and Collaboration AuroraSci. New approach to the sign problem in quantum field theories: High density qcd on a lefschetz thimble. *Physical Review D*, 86(7), 2012. ISSN 1550-7998. doi: 10.1103/PhysRevD.86.074506. URL https://doi.org/10.1103/PhysRevD.86.074506.

[64] Andrei Alexandru, Gokce Basar, Paulo F. Bedaque, Gregory W. Ridgway, and Neill C. Warrington. Sign problem and monte carlo calculations beyond lefschetz thimbles. *Journal of High Energy Physics*, 2016(5): 053, 2016. ISSN 1029-8479. doi: 10.1007/jhep05(2016)053. URL https://doi.org/10.1007/jhep05(2016)053.

[65] Marco Cristoforetti, Francesco Di Renzo, Abhishek Mukherjee, and Luigi Scorzato. Monte carlo simulations on the lefschetz thimble: Taming the sign problem. *Physical Review D*, 88(5), 2013. ISSN 1550-7998. doi: 10.1103/PhysRevD.88.051501. URL https://doi.org/10.1103/PhysRevD.88.051501.

[66] Andrei Alexandru, Paulo F. Bedaque, Henry Lamm, and Scott Lawrence. Finite-density monte carlo calculations on sign-optimized manifolds. *Physical Review D*, 97(9):094510, 2018. ISSN 2470-0010. doi: 10.1103/PhysRevD.97.094510. URL https://link.aps.org/doi/10.1103/PhysRevD.97.094510.

[67] Yuto Mori, Kouji Kashiwa, and Akira Ohnishi. Toward solving the sign problem with path optimization method. *Physical Review D*, 96

(11), 2017. ISSN 2470-0010. doi: 10.1103/PhysRevD.96.111501. URL https://doi.org/10.1103/PhysRevD.96.111501.

[68] Masafumi Fukuma, Nobuyuki Matsumoto, and Naoya Umeda. Applying the tempered lefschetz thimble method to the hubbard model away from half filling. *Physical Review D*, 100(11), 2019. ISSN 2470-0010. doi: 10.1103/PhysRevD.100.114510. URL https://doi.org/10.1103/PhysRevD.100.114510.

[69] Gert Aarts. Can stochastic quantization evade the sign problem? the relativistic bose gas at finite chemical potential. *Physical Review Letters*, 102(13):131601, 2009. doi: 10.1103/PhysRevLett.102.131601. URL https://link.aps.org/doi/10.1103/PhysRevLett.102.131601.

[70] Scott Lawrence and Yukari Yamauchi. Normalizing flows and the real-time sign problem. *Physical Review D*, 103(11), 2021. ISSN 2470-0010. doi: 10.1103/PhysRevD.103.114509. URL https://doi.org/10.1103/PhysRevD.103.114509.

[71] F. Barahona. On the computational complexity of ising spin glass models. *Journal of Physics A: Mathematical and General*, 15(10): 3241, 1982. ISSN 0305-4470. doi: 10.1088/0305-4470/15/10/028. URL https://doi.org/10.1088/0305-4470/15/10/028.

[72] Milad Marvian, Daniel A. Lidar, and Itay Hen. On the computational complexity of curing non-stoquastic hamiltonians. *Nature Communications*, 10(1):1571, 2019. doi: 10.1038/s41467-019-09501-6. URL https://doi.org/10.1038/s41467-019-09501-6.

[73] Joel Klassen, Milad Marvian, Stephen Piddock, Marios Ioannou, Itay Hen, and Barbara M. Terhal. Hardness and ease of curing the sign problem for two-local qubit hamiltonians. *SIAM Journal on Computing*, 49(6):1332–1362, 2020. doi: 10.1137/19M1287511. URL https://epubs.siam.org/doi/10.1137/19M1287511.

[74] Dominik Hangleiter, Ingo Roth, Daniel Nagaj, and Jens Eisert. Easing the monte carlo sign problem. *Science Advances*, 6(33), 2020. ISSN 2375-2548. doi: 10.1126/sciadv.abb8341. URL https://doi.org/10.1126/sciadv.abb8341.

[75] M. B. Hastings. How quantum are non-negative wavefunctions? *Journal of Mathematical Physics*, 57(1):015210, 2016. ISSN 0022-2488. doi: 10.1063/1.4936216. URL https://doi.org/10.1063/1.4936216.

[76] Z. Ringel and D. L. Kovrizhin. Quantized gravitational responses, the sign problem, and quantum complexity. *Science Advances*, 3(9):e1701758, 2017. doi: 10.1126/sciadv.1701758. URL https://www.science.org/doi/10.1126/sciadv.1701758.

[77] Omri Golan, Adam Smith, and Zohar Ringel. Intrinsic sign problem in fermionic and bosonic chiral topological matter. *Physical Review Research*, 2(4), 2020. doi: 10.1103/PhysRevResearch.2.043032. URL https://doi.org/10.1103/PhysRevResearch.2.043032.

[78] Adam Smith, Omri Golan, and Zohar Ringel. Intrinsic sign problems in topological quantum field theories. *Physical Review Research*, 2(3), 2020. doi: 10.1103/PhysRevResearch.2.033515. URL https://doi.org/10.1103/PhysRevResearch.2.033515.

[79] Donghae Seo, Minyoung You, Hee-Cheol Kim, and Gil Young Cho. Most two-dimensional bosonic topological orders forbid sign-problem-free quantum monte carlo simulation: Nonpositive gauss sum as an indicator. *Physical Review Letters*, 135(14):146602, 2025. doi: 10.1103/c8n5-l11j. URL https://link.aps.org/doi/10.1103/c8n5-l11j.

[80] Leyna Shackleton. Twisted quantum doubles are sign problem-free. *Physical Review Letters*, 136(18):186503, 2026. doi: 10.1103/j6h1-hzyz. URL https://link.aps.org/doi/10.1103/j6h1-hzyz.

[81] Rubem Mondaini, Sabyasachi Tarat, and Richard T Scalettar. Universality and critical exponents of the fermion sign problem. *arXiv preprint arXiv:2207.09026*, 2022.