Exact Phase Transition at $n \approx d^2/4$ for Centered Ellipsoid Fitting of Gaussian Points

Liz Lemma   Future Detective

January 17, 2026

Abstract

We study the feasibility of fitting n random points in $\mathbb{R}^d$ drawn i.i.d. from $\mathcal{N}(0, I_d/d)$ by the boundary of a centered ellipsoid. Writing an ellipsoid as $\{x : x^\top Sx = 1\}$ with $S \succeq 0$, this becomes a semidefinite feasibility problem with rank-one quadratic measurements $x_i^\top Sx_i$. Prior work conjectured a sharp satisfiability threshold at $n/d^2 = 1/4$ (Saunderson–Parrilo–Willsky), while recent results identify a sharp threshold at 1/4 only for bounded-spectrum approximate fits (Maillard–Bandeira, 2023). We propose an approach to upgrade the approximate bounded transition to an exact feasibility transition by proving a boundedness/regularity principle for PSD feasibility under Gaussian rank-one measurements and a 'polishing' argument that converts bounded approximate solutions into exact feasible solutions without leaving the PSD cone. The main theorem establishes that perfect ellipsoid fitting is possible w.h.p. for $n \leq (1/4 - \varepsilon)d^2$ and impossible w.h.p. for $n \geq (1/4 + \varepsilon)d^2$, and that feasible instances admit solutions with uniformly bounded operator norm. This gives a flagship instance where Gaussian-width heuristics correctly predict an exact conic satisfiability phase transition for structured measurements, with implications for SDP phase transitions and modern quadratic-feature learning models.

Table of Contents

1. Introduction and main results: statement of the exact threshold at 1/4, discussion of history (SPW13 conjecture, MB23 approximate threshold), and implications for SDP feasibility and learning theory.

2. Geometric and convex-analytic reformulation: ellipsoids as PSD matrices, linear operator $\mathcal{A}$ induced by rank-one projectors $x_i x_i^\top$, primal feasibility and dual certificates.

3. Background toolbox: conic integral geometry (Gaussian width), Gordon-type heuristics, volume universality (as in MB23), and why rank-one structure complicates exact feasibility.

4. Step I (Approximate bounded transition as a starting point): restate MB23's $\mathrm{EFP}_{\varepsilon,M}$ transition at $\alpha = 1/4$ and set up the 'approximate-to-exact' reduction target.

5. Step II (Boundedness principle for exact feasibility): prove that in the satisfiable regime, any feasibility implies existence of a bounded-spectrum feasible point; rule out 'ill-behaved only' solutions using dual certificates and conic separation.

6. Step III (Polishing bounded approximate fits to exact fits): construct a correction Δ to eliminate residual constraints while keeping PSD; develop cone-restricted invertibility for $\mathcal{A}$ and eigenvalue stability under perturbations.

7. Unsatisfiable regime: show that for $\alpha > 1/4 + \varepsilon$ no exact fit exists; reduce to nonexistence of bounded approximate fits and exclude unbounded solutions via the boundedness lemma (or by a direct dual witness).

8. Algorithmic corollaries: a polynomial-time procedure to find S (e.g., feasibility SDP / minimum-trace factor analysis); certification of infeasibility above threshold; discussion of numerical verification requirements (only for constants).

9. Extensions and robustness: universality beyond Gaussian, approximate isotropy, and perturbations; open problems (e.g., sharp finite-size scaling, geometry of the feasible set).

1 Introduction and main results: statement of the exact threshold at $1/4$, discussion of history (SPW13 conjecture, MB23 approximate threshold), and implications for SDP feasibility and learning theory.

We study the following basic question in high-dimensional geometry: given random points $x_1, \ldots, x_n \in \mathbb{R}^d$, when do they lie *exactly* on the boundary of some centered ellipsoid? Writing a centered ellipsoid as

$$E(S) := \{x \in \mathbb{R}^d : \ x^\top S x = 1\}, \qquad S \succeq 0,$$

the problem becomes the feasibility of the quadratic system $x_i^\top S x_i = 1$ for all $i \in [n]$. Throughout we take $x_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d/d)$, so that $\mathbb{E}\|x_i\|_2^2 = 1$, and we focus on the proportional regime $n = \alpha d^2$ with α held fixed as $d \to \infty$. The scale d^2 is not an artifact of the proof: the unknown $S \in \mathbb{S}^d$ has $\binom{d+1}{2}$ degrees of freedom, so the number of constraints must grow quadratically with d for a nontrivial transition to occur.

Our main conclusion is that the random feasibility event admits a sharp threshold at the constant $\alpha = 1/4$. In the satisfiable regime $\alpha < 1/4$, a fitting ellipsoid exists with high probability, whereas above $1/4$ feasibility fails with high probability. While an informal heuristic for a d^2-scale transition can be extracted from dimension counting or genericity, the appearance of the constant $1/4$ is genuinely geometric: it is the same constant that governs phase transitions for Gaussian conic feasibility problems in which the feasible set is the positive semidefinite cone $\mathbb{S}_+^d$. For truly i.i.d. Gaussian measurements (e.g. $\langle G_i, S\rangle = 1$ with G_i symmetric Gaussian), conic integral geometry predicts a transition at the squared Gaussian width of $\mathbb{S}_+^d \cap \mathbb{S}^{d(d+1)/2-1}$, which is asymptotic to $d^2/4$. Our setting is subtler: the measurements are rank-one projectors $x_i x_i^\top$, which are neither independent entries nor full-dimensional in any naive sense, and the proof must exploit both the rotational invariance of x_i and the special structure induced by rank-one observations.

The exact feasibility threshold has been explicitly conjectured in earlier work. In particular, **?** formulated the conjecture that random points admit a centered ellipsoid fit precisely up to $\alpha = 1/4$ (in the above scaling), motivated by numerical experiments and analogies with random convex programs. A major step toward this conjecture was made recently by **?**, who established a sharp transition at $\alpha = 1/4$ for a *bounded-spectrum approximate* fitting notion. Roughly speaking, they showed that when one restricts attention to matrices $S \succeq 0$ with $\|S\|_{\text{op}}$ bounded by a constant independent of d, the average constraint violation $(1/n)\sum_i |x_i^\top S x_i - 1|$ can be driven to $O(1/\sqrt{d})$ if and only if $\alpha < 1/4$. This approximate model is natural for at least two reasons. First, for $S = I_d$ we have $x_i^\top S x_i = \|x_i\|_2^2$, which concentrates around 1 with fluctuations of order $1/\sqrt{d}$, so an $O(1/\sqrt{d})$ error level is the correct baseline. Second, bounding $\|S\|_{\mathrm{op}}$ excludes extremely ill-conditioned ellipsoids, a restriction that is unavoidable in many algorithmic and statistical applications.

However, approximate feasibility under a bounded-spectrum constraint does not immediately resolve the exact problem. There are two gaps that must be closed. Below the transition, one must show that a matrix achieving small average error can be *polished* into an exact feasible solution without losing positive semidefiniteness. Above the transition, one must rule out the possibility that exact feasibility persists via pathological solutions with $\|S\|_{\mathrm{op}} \to \infty$, i.e. ellipsoids that fit all points exactly but only by becoming increasingly degenerate in some directions. Such solutions are not prohibited by MB23-type statements, since those statements explicitly impose a spectral bound. The core contribution here is to show that these gaps can in fact be closed, thereby upgrading the approximate transition at $1/4$ to an exact feasibility transition at the same location.

On the satisfiable side $\alpha < 1/4$, the key point is that bounded-spectrum approximate feasibility is not merely a surrogate for exact feasibility; it is a robust precursor. In the relevant regime, the measurement operator induced by the rank-one matrices $x_i x_i^\top$ is well-conditioned on appropriate cones of perturbations, so that one can solve a linear correction problem $\mathcal{A}(\Delta) = \mathbf{1} - \mathcal{A}(S_0)$ with Δ controlled in norm. Once the correction is small in operator norm compared to the spectral margin of S_0 (or compared to its positive eigenspace structure), deterministic eigenvalue perturbation bounds ensure that $S_0 + \Delta \succeq 0$. This "approximate-to-exact" implication is the technical mechanism by which MB23-type approximate solutions are converted into exact fits with no residual error. From a geometric viewpoint, we are exploiting that, below the transition, the intersection of the affine constraint set with $\mathbb{S}_+^d$ is not only nonempty but also stable under small perturbations of the constraints.

On the unsatisfiable side $\alpha > 1/4$, the central issue is to preclude ill-conditioned exact fits. The logical structure we use is a boundedness implication: if an exact feasible solution S exists at all, then from it we can construct a bounded-spectrum approximate solution $\widetilde{S}$ with vanishing average error as the spectral bound M grows. Intuitively, we truncate the spectrum of S (or otherwise regularize it) to obtain $\widetilde{S}$ with $\|\widetilde{S}\|_{\mathrm{op}} \leq M$, and we show that the truncation affects the quadratic forms $x_i^\top S x_i$ only mildly on average for Gaussian x_i. A nontrivial part of the argument is to quantify this "mildly" in a way that is uniform along the $n = \alpha d^2$ scaling. Once this implication is established, MB23's result applies contrapositively: if no bounded-spectrum approximate fit exists for $\alpha > 1/4$, then no exact fit can exist either, regardless of how large $\|S\|_{\mathrm{op}}$ is allowed to be. This shows that the transition at $1/4$ is not an artifact of excluding degenerate ellipsoids; rather, degeneracy cannot salvage feasibility above the threshold.

In addition to resolving the conjectured threshold, the above reasoning has consequences for semidefinite programming feasibility more broadly. The ellipsoid fitting problem is a feasibility SDP with rank-one constraints, and random instances of such SDPs arise naturally in relaxations of nonconvex problems and in randomized constructions of geometric objects. The statement that feasibility undergoes a sharp transition at a constant α is a concrete manifestation of the general phenomenon that random conic programs are governed by the intrinsic geometry of the underlying cone. Here the cone is $\mathbb{S}^d_+$, and the constant $1/4$ reflects its statistical dimension relative to the ambient space $\mathbb{S}^d$. Our contribution can thus be viewed as identifying a setting where rank-one structure does not destroy the conic-geometric prediction, and moreover where exact feasibility inherits the same transition as bounded-spectrum approximate feasibility.

There are also algorithmic and statistical implications. Below $\alpha = 1/4$, the existence of a bounded-spectrum exact fit implies that one can, with high probability, find a fitting ellipsoid in polynomial time by standard convex optimization tools, possibly supplemented with a final correction step enforcing exact constraints. Such a procedure can be interpreted as a stable method for fitting a quadratic form to data in the high-dimensional proportional regime. In learning-theoretic language, one may view S as defining a Mahalanobis-type geometry or a quadratic classifier of the form $x \mapsto x^\top Sx$. Exact fitting corresponds to placing all training points on a prescribed level set, while boundedness of $\|S\|_{\mathrm{op}}$ controls the complexity of the hypothesis class and is closely related to margin or condition-number constraints. The sharpness of the transition thus quantifies a precise sample-size barrier for such constrained quadratic representations in the random-design model.

Finally, we emphasize the conceptual message about "well-behaved" solutions. The possibility of exact feasibility via unbounded spectra would create an uncomfortable dichotomy: the primal feasibility problem would be satisfiable, but only through solutions that are unstable, poorly conditioned, and essentially invisible to bounded regularization. Our results rule out this pathology in the present model. Above the transition, infeasibility is certified in a strong sense; below it, feasibility can be achieved with uniform spectral control. In particular, the threshold at $\alpha = 1/4$ is simultaneously a threshold for existence and a threshold for existence of solutions that remain controlled as $d \to \infty$.

We next reformulate the problem in a form suitable for analysis. The quadratic constraints $x_i^\top Sx_i = 1$ can be written as linear equations in the lifted variable S against the rank-one matrices $X_i := x_i x_i^\top$, leading naturally to a linear operator $\mathcal{A} : \mathbb{S}^d \to \mathbb{R}^n$ and its adjoint $\mathcal{A}^*$. This convex-analytic viewpoint makes duality available and allows us to articulate infeasibility via explicit dual certificates, while also isolating the operator-theoretic properties needed for polishing approximate solutions into exact ones.

2 Geometric and convex-analytic reformulation: ellipsoids as PSD matrices, linear operator $\mathcal{A}$ induced by rank-one projectors $x_i x_i^\top$, primal feasibility and dual certificates.

We work throughout in the Euclidean space of real symmetric matrices $\mathbb{S}^d$, equipped with the trace inner product

$$\langle U, V \rangle := \mathrm{Tr}(UV), \qquad U, V \in \mathbb{S}^d,$$

so that $\mathbb{S}^d_+ \subset \mathbb{S}^d$ is a closed convex cone and $(\mathbb{S}^d_+)^* = \mathbb{S}^d_+$. The basic observation is that each quadratic constraint $x_i^\top S x_i = 1$ can be written linearly in S by introducing the rank-one projector

$$X_i := x_i x_i^\top \in \mathbb{S}^d_+,$$

since $x_i^\top S x_i = \mathrm{Tr}(S x_i x_i^\top) = \langle S, X_i \rangle$. This leads to a linear measurement operator $\mathcal{A} : \mathbb{S}^d \to \mathbb{R}^n$ defined by

$$(\mathcal{A}(S))_i := \langle S, X_i \rangle, \qquad i \in [n],$$

and its adjoint $\mathcal{A}^* : \mathbb{R}^n \to \mathbb{S}^d$,

$$\mathcal{A}^*(y) = \sum_{i=1}^n y_i X_i,$$

characterized by $\langle y, \mathcal{A}(S) \rangle_{\mathbb{R}^n} = \langle \mathcal{A}^*(y), S \rangle$ for all $S \in \mathbb{S}^d$ and $y \in \mathbb{R}^n$. In this notation, exact centered ellipsoid fitting becomes the conic feasibility problem

$$\text{find } S \in \mathbb{S}^d_+ \text{ such that } \mathcal{A}(S) = \mathbf{1} \in \mathbb{R}^n. \tag{1}$$

Equivalently, we ask whether the affine subspace $\{S \in \mathbb{S}^d : \mathcal{A}(S) = \mathbf{1}\}$ intersects the cone $\mathbb{S}^d_+$. The feasible set (when nonempty) is thus a spectrahedron, and many geometric questions about ellipsoid fitting can be phrased as questions about random affine slices of $\mathbb{S}^d_+$.

Two elementary remarks clarify what is gained by the lifting (1). First, convexity becomes explicit: the constraint $S \succeq 0$ is convex, and $\mathcal{A}(S) = \mathbf{1}$ is an affine condition. Second, the formulation isolates the random object that drives the phase transition, namely the linear map $\mathcal{A}$ induced by the random rank-one matrices X_i. In particular, all probabilistic statements about feasibility can be viewed as statements about the position of $\mathrm{range}(\mathcal{A}^*) = \mathrm{span}\{X_1, \ldots, X_n\} \subset \mathbb{S}^d$ relative to the cone $\mathbb{S}^d_+$, or dually about the kernel of $\mathcal{A}$. While the original variables live in $\mathbb{R}^d$, the lifted geometry takes place in ambient dimension $\dim(\mathbb{S}^d) = d(d+1)/2$, and the proportional scaling $n = \alpha d^2$ is precisely the regime in which random subspaces of $\mathbb{S}^d$ undergo nontrivial conic-intersection transitions.

The conic viewpoint also provides a canonical language for infeasibility. Since $\mathbb{S}_+^d$ is a closed convex cone and $\mathcal{A}$ is linear, standard separation principles imply a Farkas-type alternative: either (1) is feasible, or there exists a separating hyperplane witnessed by a vector $y \in \mathbb{R}^n$ such that $\mathcal{A}^*(y)$ belongs to the dual cone and the affine right-hand side $\mathbf{1}$ is strictly separated. Concretely, if we consider the (trivial-objective) primal conic program

$$\begin{array}{ll} \text{minimize} & 0 \\ \text{subject to} & \mathcal{A}(S) = \mathbf{1}, \\ & S \succeq 0, \end{array} \tag{2}$$

then its Lagrangian with multiplier $y \in \mathbb{R}^n$ is

$$\mathcal{L}(S, y) = \langle y, \mathbf{1} - \mathcal{A}(S) \rangle = \langle y, \mathbf{1} \rangle - \langle \mathcal{A}^*(y), S \rangle.$$

The associated dual program is

$$\begin{array}{ll} \text{maximize} & \langle y, \mathbf{1} \rangle \\ \text{subject to} & \mathcal{A}^*(y) \preceq 0, \end{array} \tag{3}$$

and the weak-duality inequality reads

$$\langle y, \mathbf{1} \rangle \leq \langle \mathcal{A}^*(y), S \rangle \leq 0 \qquad \text{for all primal-feasible } S \succeq 0 \text{ and dual-feasible } y \text{ with } \mathcal{A}^*(y) \preceq 0.$$

Equivalently, by flipping signs, we may state the dual feasibility condition as $\mathcal{A}^*(y) \succeq 0$ and seek $\langle y, \mathbf{1} \rangle < 0$, which is the form most convenient for certificates. Indeed, if there exists $y \in \mathbb{R}^n$ such that

$$\mathcal{A}^*(y) \succeq 0 \qquad \text{and} \qquad \langle y, \mathbf{1} \rangle < 0, \tag{4}$$

then no $S \succeq 0$ can satisfy $\mathcal{A}(S) = \mathbf{1}$, since otherwise

$$0 \leq \langle \mathcal{A}^*(y), S \rangle = \langle y, \mathcal{A}(S) \rangle = \langle y, \mathbf{1} \rangle < 0,$$

a contradiction. We refer to any y satisfying (4) as a *dual certificate* (or dual witness) of infeasibility. In our random model, constructing such a certificate amounts to choosing weights $y_1, \ldots, y_n$ so that the weighted sum $\sum_i y_i x_i x_i^\top$ is positive semidefinite, while the scalar sum $\sum_i y_i$ is strictly negative.

The structure of $\mathcal{A}^*(y)$ is worth emphasizing. Since each X_i is rank one, $\mathcal{A}^*(y)$ is a (possibly signed) weighted sample covariance matrix. If all weights y_i are nonnegative, then $\mathcal{A}^*(y) \succeq 0$ holds automatically, but then $\langle y, \mathbf{1} \rangle = \sum_i y_i \geq 0$ and (4) cannot occur. Thus any certificate must necessarily involve a nontrivial cancellation among the weights, with enough negative mass to make $\sum_i y_i < 0$ while still leaving the matrix $\sum_i y_i x_i x_i^\top$ positive semidefinite. This cancellation perspective is useful later, since it frames infeasibility as the existence of a signed reweighting of the data whose covariance becomes

PSD in a way incompatible with the imposed normalization $\mathcal{A}(S) = \mathbf{1}$. It also indicates why arguments based purely on entrywise independence are unavailable: the random matrices X_i have a rigid eigenstructure (one nonzero eigenvalue equal to $\|x_i\|_2^2$), and $\mathcal{A}^*(y)$ lives in the span of these rank-one directions.

We can also interpret (1) as a random linear system with a cone constraint. Writing $m := d(d+1)/2$, we may identify $\mathbb{S}^d$ with $\mathbb{R}^m$ via any orthonormal basis for the trace inner product, and $\mathcal{A}$ with an $n \times m$ matrix whose ith row is the vectorization of X_i. In this representation, $\mathcal{A}$ is far from an i.i.d. Gaussian matrix: the rows are dependent across coordinates because they originate from rank-one outer products. Nevertheless, $\mathbb{E}[X_i] = I_d/d$ and the distribution is rotationally invariant in the sense that for any orthogonal $Q \in \mathbb{R}^{d\times d}$, $QX_iQ^\top$ has the same law as X_i. This symmetry is the primary substitute for entrywise independence; it suggests that the relevant geometry should be invariant under conjugation and hence expressible in terms of spectral data.

There is a complementary, nonconvex but often insightful, parameterization of the primal constraint. If $S \succeq 0$ then $S = UU^\top$ for some $U \in \mathbb{R}^{d\times r}$, and the constraints become

$$x_i^\top UU^\top x_i = \|U^\top x_i\|_2^2 = 1 \qquad (i \in [n]).$$

Thus, fitting an ellipsoid is equivalent to finding a linear map $U^\top : \mathbb{R}^d \to \mathbb{R}^r$ that sends all data points to the unit sphere in $\mathbb{R}^r$. The convex lifting replaces the choice of U (and the unknown rank r) by the matrix S, in exchange for enlarging the ambient dimension but gaining convexity. This viewpoint helps explain why the operator norm $\|S\|_{\mathrm{op}}$ plays a privileged role: $\|S\|_{\mathrm{op}} = \|U\|_{\mathrm{op}}^2$ controls the Lipschitz constant of the embedding $x \mapsto U^\top x$, and hence bounds how violently the map can distort directions in $\mathbb{R}^d$. Pathological exact fits above the transition would correspond to maps U with exploding operator norm that place the points on a sphere only by collapsing or stretching certain directions extremely.

Both the primal and dual formulations highlight a general dichotomy: below the transition we seek to establish existence of an S with quantitative spectral control, whereas above the transition we seek either a dual certificate y as in (4) or an argument showing that any primal solution would force a well-behaved approximate solution. In either route, the probabilistic challenge is to understand the random operator $\mathcal{A}$ induced by $X_i = x_i x_i^\top$. In particular, one repeatedly encounters two competing effects. On the one hand, because each x_i is Gaussian and rotationally invariant, linear and quadratic forms in x_i enjoy strong concentration and admit comparison principles with genuinely Gaussian objects. On the other hand, the rank-one structure imposes heavy algebraic constraints: the measurements $\langle S, X_i\rangle$ are not independent across different S, and $\mathcal{A}$ is not an isotropic embedding of $\mathbb{S}^d$ in the same sense as an i.i.d. Gaussian measurement ensemble.

To prepare for the quantitative tools that follow, it is useful to isolate the geometric primitives that enter the analysis. The cone $\mathbb{S}_+^d$ is self-dual and invariant under conjugation by orthogonal matrices; its intrinsic geometry can be captured by notions such as statistical dimension and Gaussian width. The measurement map $\mathcal{A}$ defines a random affine slice of this cone, and feasibility depends on whether this slice hits the cone in a region that is not too close to the boundary (in the satisfiable regime) or, conversely, whether the slice can be strictly separated from the cone (in the unsatisfiable regime). When we speak of "polishing" an approximate solution, we are implicitly using that $\mathcal{A}$ is well-conditioned on certain tangent or descent cones determined by the PSD constraint and the spectral regularity of the approximate solution. When we speak of a dual witness, we are using that a random span of rank-one projectors typically intersects the dual cone in a way that permits strict separation once the number of constraints exceeds the critical value.

At this point, the problem has been reduced to an interplay between (i) conic geometry of $\mathbb{S}_+^d$ in ambient dimension $\Theta(d^2)$ and (ii) random linear maps generated by rank-one Gaussian projectors. The remaining task is to import a toolbox that can make this interplay quantitative at the level of constants. In the i.i.d. Gaussian measurement model, sharp thresholds are governed by conic integral geometry and comparison inequalities of Gordon type; in our rank-one model, we require universality inputs and careful replacements for isotropy and independence. We therefore turn next to the relevant background: Gaussian widths and statistical dimension heuristics, Gordon-type arguments, and the volume-universality phenomena leveraged by **?**, along with an explanation of why the rank-one structure is precisely the point at which the exact feasibility analysis becomes nontrivial.

3 Background toolbox: conic integral geometry (Gaussian width), Gordon-type heuristics, volume universality (as in MB23), and why rank-one structure complicates exact feasibility.

We now summarize the probabilistic and geometric toolbox that governs threshold phenomena for conic feasibility, and we explain why the rank-one structure $X_i = x_i x_i^\top$ forces additional work when we pass from bounded-spectrum approximate statements to exact feasibility. The guiding principle is that in ambient dimension $m = \dim(\mathbb{S}^d) = d(d+1)/2$, the event that a random affine slice intersects a closed convex cone is controlled, to first order, by a single scalar parameter of the cone (its *statistical dimension*), and that the constant $1/4$ arises from the fact that $\delta(\mathbb{S}_+^d) \sim d^2/4$. The technical burden is to justify that our rank-one measurement operator $\mathcal{A}$ exhibits the same effective geometry as an i.i.d. Gaussian operator on the relevant sets of matrices.

Gaussian width and statistical dimension. Let $g \sim \mathcal{N}(0, I_m)$ and let $T \subset \mathbb{R}^m$ be bounded. The *Gaussian width* of T is

$$w(T) := \mathbb{E}\Big[\sup_{t \in T} \langle g, t \rangle\Big].$$

For a closed convex cone $C \subset \mathbb{R}^m$, one typically takes $T = C \cap \mathbb{S}^{m-1}$, so that $w(C \cap \mathbb{S}^{m-1})$ measures the "size" of the cone as seen by a Gaussian direction. Closely related is the *statistical dimension*

$$\delta(C) := \mathbb{E}\big[\|\Pi_C(g)\|_2^2\big],$$

where Π_C denotes Euclidean projection onto C. For cones, these quantities satisfy inequalities of the form

$$w(C \cap \mathbb{S}^{m-1})^2 \;\leq\; \delta(C) \;\leq\; w(C \cap \mathbb{S}^{m-1})^2 + 1,$$

so either can be used as the governing parameter at the level of leading constants. The statistical dimension enjoys particularly clean identities and kinematic formulas; it is also stable under orthogonal transformations, which is essential for rotationally invariant models.

In our setting the ambient space is $(\mathbb{S}^d, \langle \cdot, \cdot \rangle)$, which we may identify with $\mathbb{R}^m$ using any orthonormal basis for the trace inner product. The cone $\mathbb{S}^d_+$ is orthogonally invariant under conjugation, and its statistical dimension is known explicitly (see, e.g., **?**): one has

$$\delta(\mathbb{S}^d_+) = \frac{d(d+1)}{4}\,, \tag{5}$$

so $\delta(\mathbb{S}^d_+)/d^2 \to 1/4$. This identity is the quantitative origin of the constant $1/4$ appearing in the phase transition.

Conic kinematics and threshold heuristics. A basic template is the following: let $C \subset \mathbb{R}^m$ be a closed convex cone and let $L \subset \mathbb{R}^m$ be a random subspace of codimension n, distributed uniformly over the Grassmannian. Conic integral geometry predicts a sharp transition for the intersection event $C \cap L \neq \{0\}$ as n passes $\delta(C)$. More precisely, the approximate kinematic formula (again, see **?**) asserts that if L is random and independent of C, then

$$\mathbb{P}[C \cap L \neq \{0\}] \approx \begin{cases} 1, & n < \delta(C), \\ 0, & n > \delta(C), \end{cases}$$

with a window of width $O(\sqrt{m})$ around $\delta(C)$. While our feasibility problem is an affine intersection rather than a homogeneous one, the same geometry appears after linearization at a point and after passing to tangent/descent cones: the local obstruction to feasibility (or the local well-conditioning needed for "polishing") is governed by whether the kernel of the measurement operator intersects an appropriate cone nontrivially.

If, hypothetically, $\mathcal{A}$ were an i.i.d. Gaussian operator on $\mathbb{S}^d$ (i.e. its matrix representation had independent $\mathcal{N}(0,1)$ entries after choosing an orthonormal basis), then $\ker(\mathcal{A})$ would be a uniformly random subspace of dimension $m-n$. The conic kinematic heuristic would then suggest that phenomena tied to the PSD constraint should undergo transitions near $n \approx \delta(\mathbb{S}^d_+) \approx d^2/4$, that is, at $\alpha = n/d^2 \approx 1/4$. This is the "cone-only" prediction: it ignores rank-one structure and uses only the orthogonal invariance of a Gaussian embedding.

Gordon-type comparisons (escape through a mesh). A complementary viewpoint is provided by Gordon's inequality for Gaussian processes, which yields quantitative lower bounds for $\inf_{t\in T} \|Gt\|_2$ when G is a Gaussian matrix. In one convenient form (escape through a mesh), if $G \in \mathbb{R}^{n\times m}$ has i.i.d. $\mathcal{N}(0,1)$ entries and $T \subset \mathbb{S}^{m-1}$ is measurable, then with high probability

$$\inf_{t\in T} \|Gt\|_2 \;\gtrsim\; \sqrt{n} - w(T),$$

up to lower-order terms. Applied with $T = C \cap \mathbb{S}^{m-1}$, this gives that G is bounded below on the cone (equivalently, $\ker(G) \cap C = \{0\}$) once $n > w(C \cap \mathbb{S}^{m-1})^2 \approx \delta(C)$. In conic optimization, such lower bounds translate into stability (restricted invertibility) on tangent or descent cones, which is exactly the type of estimate one needs for a polishing step: if the residual lies in the range of $\mathcal{A}$ and $\mathcal{A}$ is well-conditioned on the relevant cone of allowable corrections, then one can solve $\mathcal{A}(\Delta) = r$ with Δ controlled in Frobenius and operator norms.

In the i.i.d. Gaussian measurement model, the Gordon mechanism essentially explains both sides of a sharp transition: above the threshold, random subspaces avoid the cone (yielding separation/infeasibility), while below the threshold, the operator is sufficiently well-conditioned on the appropriate cones to allow correction and stability. Thus, if we could replace our rank-one $\mathcal{A}$ by a Gaussian G without changing the relevant conic geometry, the constant $1/4$ would be immediate from (5).

What "universality" must accomplish. Our operator $\mathcal{A}$ is not Gaussian in $\mathbb{S}^d$: its ith measurement is the quadratic form $\langle S, X_i\rangle = x_i^\top S x_i$, and as S varies the family $\{x_i^\top S x_i\}$ is a second-order Gaussian chaos rather than a linear Gaussian process with independent coordinates. Moreover, the rows of the matrix representation of $\mathcal{A}$ (the vectorizations of X_i) are not isotropic in $\mathbb{R}^m$ without a nontrivial covariance correction, and they are constrained to lie in the rank-one manifold inside $\mathbb{S}^d$. Thus, to import conic integral geometry heuristics we need a *universality principle*: at the scale $n = \Theta(d^2)$, the random affine slice induced by $\mathcal{A}$ should behave like a slice induced by a rotationally invariant ensemble in $\mathbb{S}^d$, at least when tested against suitably regular sets of matrices.

The relevant notion of universality here is not entrywise universality (which would fail because the vectorized X_i are highly dependent across coordinates), but rather *volume* or *conic* universality: random convex bodies or random cones generated by the measurements should have asymptotically the same intrinsic volumes as those generated by a Gaussian ensemble. Informally, when we restrict attention to sets of matrices with bounded operator norm (or bounded condition number), the anisotropy introduced by the rank-one structure averages out by rotational invariance of the x_i, and the resulting geometry becomes indistinguishable from the Gaussian benchmark.

Volume universality in the sense of ?. The work **?** establishes precisely such a statement for an *approximate* version of ellipsoid fitting with explicit spectral regularization. While we defer the formal restatement to the next step, we emphasize here what is conceptually gained. First, by imposing $\|S\|_{\text{op}} \le M$ and allowing average constraint violation on the order $\varepsilon/\sqrt{d}$, one restricts attention to a compact, well-behaved subset of $\mathbb{S}^d_+$ on which concentration and metric-entropy arguments can be made uniform in d. Second, the approximate constraints naturally match the scale of fluctuations of $\|x_i\|_2^2$ and, more generally, of $x_i^\top S x_i$ when $\mathrm{Tr}(S)$ is $O(d)$ and $\|S\|_{\text{op}}$ is $O(1)$: indeed, for such S, one has $\mathbb{E}[x_i^\top S x_i] = \mathrm{Tr}(S)/d$ and $\mathrm{Var}(x_i^\top S x_i) = \Theta(\|S\|_F^2/d^2)$, which is typically $\Theta(1/d)$ when $\|S\|_F = \Theta(\sqrt{d})$. Thus the normalization $\varepsilon/\sqrt{d}$ reflects the intrinsic noise level of the rank-one measurements.

At a high level, volume-universality arguments proceed by comparing the random image (or preimage) of a fixed convex set under two ensembles: the rank-one ensemble induced by $x_i x_i^\top$ and a Gaussian ensemble on $\mathbb{S}^d$. The comparison is not pointwise; instead it controls global quantities such as intrinsic volumes, support functions after smoothing, or probabilities of separation for families of halfspaces. The bounded-spectrum condition is essential: it ensures that the relevant processes (suprema over S in a bounded set) are governed by Lipschitz functionals of the data with respect to a norm compatible with Gaussian comparison principles. In turn, this allows one to show that the approximate feasibility event has the same limiting behavior as in the i.i.d. Gaussian measurement model, hence exhibiting a transition at the constant predicted by (5), namely 1/4.

Why rank-one structure complicates exact feasibility. Exact feasibility removes both regularizers at once: we require $\mathcal{A}(S) = \mathbf{1}$ exactly, and we do not a priori bound $\|S\|_{\text{op}}$. This creates two distinct obstructions.

First, the rank-one structure makes $\mathcal{A}$ potentially poorly conditioned on directions associated with high condition number matrices. If S has a very large top eigenvalue, then the random variables $x_i^\top S x_i$ become sensitive to rare alignments of x_i with the top eigenvector of S. In a Gaussian measurement model on $\mathbb{S}^d$ such rare alignments are averaged over independent coordinates; in the rank-one model they interact with the single-direction nature of each X_i. Consequently, uniform concentration over unbounded families of S is unavailable, and it is not legitimate to directly pass from the approximate transition to the exact one without showing that any exact solution can be regularized.

Second, exact feasibility is a boundary question for the cone: feasible spectrahedra can in principle touch $\partial\mathbb{S}_+^d$ in complicated ways, and the existence of a feasible point does not automatically yield a feasible point with a spectral gap. For polishing arguments, however, one needs a quantitative interiority statement (or at least control on the smallest nonzero eigenvalues on the support) in order to perturb while preserving positive semidefiniteness. In the Gaussian measurement model, conic integral geometry can often be invoked directly to show that random slices intersect the cone transversely, yielding typical nondegeneracy. In the rank-one model, transversality must be proved by hand on the relevant cones, and it can fail if one allows S to approach extreme ill-conditioning.

These obstructions explain the division of labor in our approach. On the satisfiable side, we will rely on bounded-spectrum approximate feasibility (a robust, "interior" statement) and then prove that $\mathcal{A}$ is well-conditioned on a cone of allowable corrections, enabling an approximate-to-exact polishing step with controlled operator norm. On the unsatisfiable side, we must rule out the possibility that feasibility persists only through ill-conditioned solutions: one must show that any putative exact feasible S can be truncated, rescaled, or otherwise regularized to produce a bounded-spectrum approximate feasible solution, contradicting the approximate infeasibility above the transition.

Analytic substitutes for i.i.d. Gaussian measurements. To execute these steps, we repeatedly use tools that play the role of Gordon-type bounds and RIP-type estimates, but adapted to quadratic measurements. For fixed S, the random variable $x^\top S x$ admits sharp tail bounds via Hanson–Wright inequalities; for families of S, one uses symmetrization and chaining for second-order Gaussian chaos, with metrics induced by $\|\cdot\|_F$ and $\|\cdot\|_{\mathrm{op}}$. At the operator level, $\mathcal{A}^*(y) = \sum_i y_i x_i x_i^\top$ is a weighted sample covariance matrix, for which nonasymptotic spectral bounds follow from matrix Bernstein-type inequalities when the weights and norms are controlled. The point is not that these tools reproduce entrywise independence, but that they allow us to prove *cone-restricted invertibility* and *spectral stability* on the specific subsets of $\mathbb{S}^d$ that arise from bounded-spectrum approximate solutions and their tangent cones.

Taken together, conic integral geometry provides the constant 1/4 and the correct qualitative picture; Gordon-type arguments indicate what quantitative conditioning statements suffice for polishing and separation; and the volume-universality results of **?** supply a rigorous starting point in the rank-one model, but only under bounded-spectrum approximate notions. The remaining work is therefore structural: we must (i) show that below the transition, bounded-spectrum approximate fits can be polished to exact fits using restricted invertibility of $\mathcal{A}$ on a suitable cone, and (ii) show that above the transition, any exact fit would necessarily induce a bounded-spectrum approximate fit, eliminating the possibility of purely ill-behaved solutions. We now begin with the approximate transition statement and formulate the reduction targets precisely.

Step I: bounded-spectrum approximate feasibility as the starting point. We introduce an approximate and regularized notion of ellipsoid fitting, which will serve as the input to all subsequent reductions. Fix an approximation tolerance parameter $\epsilon > 0$ and an operator-norm bound $M \in (0, \infty)$. We consider the event

$$\mathrm{EFP}_{\epsilon,M}(n,d): \quad \exists S \in \mathbb{S}_+^d \ \text{ s.t. } \ \|S\|_{\mathrm{op}} \le M \ \text{ and } \ \frac{1}{n}\sum_{i=1}^n \left| x_i^\top S x_i - 1 \right| \le \frac{\epsilon}{\sqrt{d}} .$$

This is the "bounded-spectrum approximate fitting" event: the quadratic constraints are enforced only on average and at the intrinsic fluctuation scale $1/\sqrt{d}$, while the spectrum of S is confined to a compact set. We stress that the constraint $\|S\|_{\mathrm{op}} \le M$ is not a technical convenience but the structural hypothesis under which volume-universality methods can be made uniform over d; without it, one must contend with genuinely non-uniform behavior driven by rare alignments of x_i with extreme eigendirections.

A useful companion formulation is an ℓ_2-residual variant, aligned with the linear correction step used in polishing. We define

$$\mathrm{EFP}^{(2)}_{\epsilon,M}(n,d): \quad \exists S \in \mathbb{S}_+^d \ \text{ s.t. } \ \|S\|_{\mathrm{op}} \le M \ \text{ and } \ \|\mathcal{A}(S) - \mathbf{1}\|_2 \le \epsilon\sqrt{\frac{n}{d}} .$$

While $\mathrm{EFP}_{\epsilon,M}$ is stated in average absolute error, it is convenient to pass between ℓ_1 and ℓ_2 residuals under bounded spectrum: for each fixed M, the random variables $x_i^\top S x_i$ are uniformly sub-exponential over the class $\{S \succeq 0 : \|S\|_{\mathrm{op}} \le M\}$, so with high probability one has a mild (e.g. polylogarithmic) control on $\max_i |x_i^\top S x_i|$; in that regime, an ℓ_1 bound of order $\epsilon/\sqrt{d}$ implies an ℓ_2 bound of order $\epsilon\sqrt{n/d}$ up to constants. Accordingly, we will treat $\mathrm{EFP}_{\epsilon,M}$ and $\mathrm{EFP}^{(2)}_{\epsilon,M}$ as interchangeable inputs, at the expense of adjusting constants depending on (ϵ, M).

MB23 transition at $\alpha = 1/4$. We now record the approximate phase transition proved in **?**, specialized to our notation $n = \alpha d^2$ and to the Gaussian design $x_i \sim \mathcal{N}(0, I_d/d)$. The statement we require is that, once $\epsilon > 0$ is fixed, bounded-spectrum approximate feasibility undergoes a sharp transition at $\alpha = 1/4$, matching the statistical-dimension prediction for $\mathbb{S}^d_+$.

Theorem (MB23, bounded-spectrum approximate transition). *Fix $\varepsilon \in (0, 1/10)$. There exist constants $\epsilon_0 = \epsilon_0(\varepsilon) > 0$ and $M_0 = M_0(\varepsilon) < \infty$ such that the following holds as $d \to \infty$ with $\alpha = n/d^2$ converging.*

1. *(Approximate satisfiability below 1/4.) If $\alpha \le \frac{1}{4} - \varepsilon$, then*
$$\mathbb{P}\big[\mathrm{EFP}_{\epsilon_0, M_0}(n, d)\big] \to 1\,.$$

2. *(Approximate unsatisfiability above 1/4.) If $\alpha \ge \frac{1}{4} + \varepsilon$, then for every fixed $M < \infty$,*
$$\mathbb{P}\big[\mathrm{EFP}_{\epsilon_0, M}(n, d)\big] \to 0\,.$$

Equivalently (up to constant renormalization), the same conclusions hold with $\mathrm{EFP}^{(2)}_{\epsilon_0,M}$ in place of $\mathrm{EFP}_{\epsilon_0,M}$.

For our purposes, the essential content is not the precise form of the tolerance (whether average absolute error or ℓ_2-residual) but the existence of *some* tolerance level $\epsilon_0(\varepsilon)$ and some boundedness level $M_0(\varepsilon)$ for which the probability transition at $\alpha = 1/4$ is rigorous in the rank-one model. In particular, this theorem identifies the same constant $1/4$ that conic integral geometry predicts for rotationally invariant Gaussian embeddings, thereby confirming that the rank-one ensemble is "conically universal" at the bounded-spectrum level.

Normalization and scale of the approximation error. It is worth making explicit why the scale $\epsilon/\sqrt{d}$ is the natural one. For any fixed $S \succeq 0$ with $\|S\|_{\mathrm{op}} \le M$ and $\mathrm{Tr}(S) = \Theta(d)$, we have

$$\mathbb{E}[x^\top S x] = \frac{\mathrm{Tr}(S)}{d} = \Theta(1)\,, \qquad \mathrm{Var}(x^\top S x) = \frac{2\|S\|_F^2}{d^2} = \Theta\left(\frac{1}{d}\right) \quad \text{when } \|S\|_F = \Theta(\sqrt{d})\,,$$

so an individual constraint $x_i^\top S x_i \approx 1$ cannot typically be enforced to better than $O(1/\sqrt{d})$ accuracy uniformly over i without overfitting. The MB23 formulation allows exactly this intrinsic scale of fluctuation while controlling the class of candidate matrices by $\|S\|_{\mathrm{op}} \le M$. Consequently, the bounded-spectrum approximate event is robust: it is insensitive to a small number of atypical x_i, and it is stable under small perturbations of S, both of which are necessary for universality and for subsequent correction arguments.

Reduction targets: from approximate to exact, and from exact to approximate. The MB23 theorem does not, by itself, imply a statement about exact feasibility $\mathrm{EFP}_{0,\infty}$, because exact feasibility removes both regularizers simultaneously: we demand $\mathcal{A}(S) = \mathbf{1}$ exactly, and we allow $\|S\|_{\mathrm{op}}$ to be arbitrarily large. Our strategy is therefore to use MB23 as a black box and reduce the exact phase transition to two additional implications, each of which is proved by a separate argument tailored to rank-one measurements.

We formalize these implications as follows. First, in the satisfiable regime $\alpha \leq \frac{1}{4} - \varepsilon$, MB23 provides (with high probability) a matrix $S_0 \succeq 0$ with bounded operator norm and small residual $r := \mathcal{A}(S_0) - \mathbf{1}$. The *approximate-to-exact* target is to solve the linear correction equation

$$\mathcal{A}(\Delta) = -r$$

with Δ controlled in operator norm so that $S_0 + \Delta \succeq 0$. In other words, we aim to show that, below the transition, the operator $\mathcal{A}$ is well-conditioned on a cone of allowable perturbations around a bounded-spectrum point, and that this conditioning suffices to preserve positive semidefiniteness after correction. This is the polishing mechanism encoded in our Theorem C.

Second, in the unsatisfiable regime $\alpha \geq \frac{1}{4} + \varepsilon$, MB23 asserts that no bounded-spectrum approximate solution exists (for the relevant tolerances). The *exact-to-approximate* target is to show that if an exact solution $S \succeq 0$ with $\mathcal{A}(S) = \mathbf{1}$ were to exist at all, then one could regularize it (by spectral truncation, rescaling, and a controlled repair of the induced residual) into a bounded-spectrum matrix $\widetilde{S}$ satisfying $\mathrm{EFP}_{\epsilon_0,M}$ for some finite M. This would contradict the MB23 impossibility statement above the transition, thereby ruling out even purely ill-conditioned exact fits. This boundedness principle is the content of our Theorem B (and, in a more elementary form, Lemma 4).

A convenient reformulation in terms of residual vectors. To keep the two reductions conceptually separate, we record the data that MB23 supplies in a form compatible with both directions. Suppose $\mathrm{EFP}^{(2)}_{\epsilon,M}(n,d)$ holds, and let S_0 be a witness. The residual vector $r = \mathcal{A}(S_0) - \mathbf{1} \in \mathbb{R}^n$ then satisfies

$$\|r\|_2 \leq \epsilon\sqrt{\frac{n}{d}}, \qquad \text{and} \qquad S_0 \succeq 0, \quad \|S_0\|_{\mathrm{op}} \leq M.$$

The polishing step seeks Δ with $\mathcal{A}(\Delta) = -r$ and $\|\Delta\|_{\mathrm{op}}$ small compared to the spectral gap on the support of S_0. Achieving this requires two analytic inputs: (i) restricted invertibility of $\mathcal{A}$ on a cone capturing admissible PSD-preserving directions, giving Frobenius control of Δ; and (ii) a mechanism upgrading Frobenius control to operator-norm control, so that Weyl-type perturbation bounds ensure $S_0 + \Delta \succeq 0$. These are precisely the roles played by our Lemmas 1–3.

Conversely, suppose $\mathrm{EFP}_{0,\infty}(n,d)$ holds, and let S be an exact witness. Since no a priori bound on $\|S\|_{\mathrm{op}}$ is available, we cannot directly apply MB23-type compactness arguments. The regularization step therefore begins by truncating the spectrum of S at a level M, producing (after possible rescaling) a matrix $\widetilde{S}$ with $\|\widetilde{S}\|_{\mathrm{op}} \le M$. The truncation introduces an error vector $\widetilde{r} := \mathcal{A}(\widetilde{S}) - \mathbf{1}$, and the task is to show that $\widetilde{r}$ is small in an averaged sense once M is chosen large enough. One then either accepts $\widetilde{S}$ as an approximate witness, or performs a further correction (now within a bounded class) to meet the precise MB23 error criterion. This is the mechanism formalized in Lemma 4 and upgraded in Theorem B.

Logical structure of the proof once MB23 is granted. With the MB23 theorem in hand, the remaining arguments are purely implications between events, proved with high probability under the same scaling $n = \alpha d^2$. Concretely, fixing $\varepsilon > 0$, we will establish:

- for $\alpha \le \frac{1}{4} - \varepsilon$: $\mathrm{EFP}_{\epsilon_0,M_0} \Rightarrow \mathrm{EFP}_{0,\infty}$ with high probability (polishing);
- for $\alpha \ge \frac{1}{4} + \varepsilon$: $\mathrm{EFP}_{0,\infty} \Rightarrow \mathrm{EFP}_{\epsilon_0,M}$ for some finite M with high probability (boundedness principle).

Combining these with the MB23 transition yields the exact threshold in Theorem A. We emphasize that neither implication is a formal consequence of conic duality alone: both use specific analytic properties of the rank-one measurement operator $\mathcal{A}$ (Hanson–Wright control for quadratic forms, chaining for Gaussian chaos on bounded classes, and matrix concentration for $\mathcal{A}^*(y)$). The key point, however, is that these analytic inputs are invoked only on bounded-spectrum sets (for polishing) or to *create* a bounded-spectrum object from an unbounded one (for boundedness). Thus MB23 serves as the sole geometric input needed to locate the constant $1/4$, and all remaining work is structural.

Transition from Step I to Step II. We now proceed to the first of the two reductions needed to upgrade the approximate transition to the exact one. Namely, we will show that exact feasibility cannot persist above $\alpha = 1/4$ by hiding in ill-conditioned matrices: any exact feasible point necessarily induces a bounded-spectrum approximate feasible point. This boundedness principle, together with a dual-separation viewpoint, is the content of Step II.

Step II: a boundedness principle and dual separation. In order to upgrade the MB23 transition for $\mathrm{EFP}_{\epsilon,M}$ to the exact event $\mathrm{EFP}_{0,\infty}$, we must exclude the following pathological possibility: for $\alpha > \frac{1}{4}$, exact feasibility could persist but only through witnesses $S \succeq 0$ with increasingly ill-conditioned spectra, thereby evading the compactness inherent in

$\|S\|_{\mathrm{op}} \le M$. Our goal in this step is to show that such "ill-behaved only" feasibility cannot occur. Concretely, we establish an implication of the form

$$\mathrm{EFP}_{0,\infty}(n,d) \implies \mathrm{EFP}_{\epsilon_0,M}(n,d)$$

for some finite $M = M(\varepsilon)$ and the tolerance $\epsilon_0(\varepsilon)$ appearing in MB23. Combined with the MB23 impossibility statement above $\alpha = \frac{1}{4}$, this yields exact infeasibility even when $\|S\|_{\mathrm{op}}$ is unconstrained.

Trace normalization forced by the constraints. We begin with a simple but crucial normalization consequence of exact feasibility. Suppose $S \succeq 0$ satisfies $\mathcal{A}(S) = \mathbf{1}$. Let

$$\widehat{\Sigma} := \frac{1}{n}\sum_{i=1}^{n} X_i = \frac{1}{n}\sum_{i=1}^{n} x_i x_i^\top.$$

Then

$$1 = \frac{1}{n}\sum_{i=1}^{n} x_i^\top S x_i = \frac{1}{n}\sum_{i=1}^{n} \mathrm{Tr}(SX_i) = \mathrm{Tr}(S\widehat{\Sigma}).$$

By standard Gaussian covariance concentration in operator norm (here $n = \alpha d^2 \gg d$), we have with high probability

$$\big\|\widehat{\Sigma} - \tfrac{1}{d}I_d\big\|_{\mathrm{op}} \le C\sqrt{\frac{d}{n}} = O\Big(\frac{1}{\sqrt{d}}\Big).$$

Using $\mathrm{Tr}(S\widehat{\Sigma}) = \mathrm{Tr}(S)/d + \mathrm{Tr}\big(S(\widehat{\Sigma} - \frac{1}{d}I)\big)$ and $|\mathrm{Tr}(S(\widehat{\Sigma} - \frac{1}{d}I))| \le \|S\|_{\mathrm{op}}\|\widehat{\Sigma} - \frac{1}{d}I\|_{\mathrm{op}}\mathrm{rank}(S)$, we obtain that any exact witness necessarily has trace of order d. In particular, on the high-probability event $\|\widehat{\Sigma} - \frac{1}{d}I\|_{\mathrm{op}} \le c_0$ with $c_0 \ll 1$, we have the deterministic implication

$$(1-c_0)\frac{\mathrm{Tr}(S)}{d} \le 1 \le (1+c_0)\frac{\mathrm{Tr}(S)}{d}, \qquad \text{hence} \qquad \mathrm{Tr}(S) = d\cdot(1+O(c_0)). \tag{6}$$

Thus an exact witness cannot concentrate all of its mass in a vanishing-dimensional subspace without generating extremely large eigenvalues; conversely, if large eigenvalues exist, they must be sparse in the sense of trace accounting.

Spectral truncation and the induced residual. Given any $S \succeq 0$ with $\mathcal{A}(S) = \mathbf{1}$, we define its spectral truncation at level M by

$$S^{(M)} := \Pi_M(S) := U\,\mathrm{diag}(\min\{\lambda_j, M\})\,U^\top,$$

where $S = U\,\mathrm{diag}(\lambda_j)\,U^\top$ is an eigen-decomposition. Set $T^{(M)} := S - S^{(M)} \succeq 0$. Then $\|S^{(M)}\|_{\mathrm{op}} \le M$ and, for each i,

$$x_i^\top S^{(M)} x_i = x_i^\top S x_i - x_i^\top T^{(M)} x_i = 1 - x_i^\top T^{(M)} x_i,$$

so the residual vector associated to $S^{(M)}$ is entrywise nonpositive:

$$r^{(M)} \;:=\; \mathcal{A}(S^{(M)}) - \mathbf{1} \;=\; -\mathcal{A}(T^{(M)}), \qquad \text{hence} \qquad |r_i^{(M)}| = x_i^\top T^{(M)} x_i.$$

The average ℓ_1 residual is therefore

$$\frac{1}{n}\sum_{i=1}^n |r_i^{(M)}| \;=\; \frac{1}{n}\sum_{i=1}^n x_i^\top T^{(M)} x_i \;=\; \operatorname{Tr}\big(T^{(M)}\widehat{\Sigma}\big). \tag{7}$$

On the event $\|\widehat{\Sigma} - \frac{1}{d}I\|_{\mathrm{op}} \le c_0$, we obtain the deterministic bound

$$\operatorname{Tr}\big(T^{(M)}\widehat{\Sigma}\big) \;\le\; \frac{1+c_0}{d}\operatorname{Tr}(T^{(M)}). \tag{8}$$

Thus, to obtain an ℓ_1 residual of order $1/\sqrt{d}$, it suffices to show that any exact witness S has $\operatorname{Tr}(T^{(M)}) = O(\sqrt{d})$ for some fixed M. This is the substantive content of the boundedness principle: exact feasibility forces the spectral tail above a fixed truncation level to be small in trace, uniformly over d.

Boundedness via a repairable truncation: the core lemma. We isolate the implication we will use in the unsatisfiable regime.

Lemma (bounded-spectrum approximate witness from an exact witness). *Fix $\varepsilon \in (0, 1/10)$. There exist constants $M_\star = M_\star(\varepsilon) < \infty$ and $c_\star = c_\star(\varepsilon) > 0$ such that the following holds with probability $1 - o(1)$ as $d \to \infty$. If there exists $S \succeq 0$ with $\mathcal{A}(S) = \mathbf{1}$, then there exists $\widetilde{S} \succeq 0$ with $\|\widetilde{S}\|_{\mathrm{op}} \le M_\star$ and*

$$\frac{1}{n}\sum_{i=1}^n \big|x_i^\top \widetilde{S} x_i - 1\big| \;\le\; \frac{c_\star}{\sqrt{d}}\,.$$

In particular, for $\epsilon_0(\varepsilon)$ as in MB23, choosing $c_\star \le \epsilon_0$ yields $\mathrm{EFP}_{0,\infty} \Rightarrow \mathrm{EFP}_{\epsilon_0, M_\star}$ with high probability.

The construction of $\widetilde{S}$ is conceptually simple: we truncate S to obtain a bounded-spectrum matrix $S^{(M)}$, then we allow a further bounded adjustment within the compact set $\{0 \preceq S \preceq MI\}$ to reduce the residual to the intrinsic $1/\sqrt{d}$ scale. The novelty is that we do *not* need to solve the correction equation $\mathcal{A}(\Delta) = -r^{(M)}$ exactly (as we will in Step III); rather, an approximate repair suffices, and can be certified by conic duality.

At a high level, we argue as follows. First, from (6), any exact witness must satisfy $\operatorname{Tr}(S) = d(1 + o(1))$. Second, we choose M large (but fixed) and show that, on a high-probability event depending only on the random design, the truncated tail $T^{(M)} = (S - MI)_+$ cannot have trace larger than $O(\sqrt{d})$; otherwise, the family of constraints $\{x_i^\top S x_i = 1\}_{i \le n}$ would force an atypical flattening of the projected norms of x_i in the top-eigenspaces of S, which can be ruled out by uniform concentration of χ^2-type statistics over subspaces of dimension up to $O(d)$. Plugging $\mathrm{Tr}(T^{(M)}) \lesssim \sqrt{d}$ into (8) yields

$$\frac{1}{n}\sum_{i=1}^{n}\left|x_i^\top S^{(M)}x_i - 1\right| \;=\; \mathrm{Tr}\big(T^{(M)}\widehat{\Sigma}\big) \;\lesssim\; \frac{1}{\sqrt{d}}.$$

Finally, if the truncation produces a residual with the correct scaling but the wrong constant, we perform a bounded repair by minimizing ℓ_1 residual over the compact convex set $\{S \succeq 0 : \|S\|_{\mathrm{op}} \le M\}$; duality ensures that, if this optimum were bounded away from 0 at the $1/\sqrt{d}$ scale, one could separate $\mathbf{1}$ from $\mathcal{A}(\{S \succeq 0 : \|S\|_{\mathrm{op}} \le M\})$ by a dual witness, contradicting the assumed exact feasibility. This "duality-based repair" is robust and does not require the restricted invertibility technology of Step III.

Conic separation viewpoint and dual certificates. We now formalize the separating-hyperplane mechanism that underlies both the boundedness principle and the eventual infeasibility result. Consider the primal feasibility problem

$$\text{find} \quad S \succeq 0 \quad \text{such that} \quad \mathcal{A}(S) = \mathbf{1}.$$

A standard conic Farkas lemma implies: exactly one of the following holds:

1. there exists $S \succeq 0$ with $\mathcal{A}(S) = \mathbf{1}$;
2. there exists $y \in \mathbb{R}^n$ such that $\mathcal{A}^*(y) \succeq 0$ and $\langle y, \mathbf{1}\rangle < 0$.

Indeed, if $\mathcal{A}^*(y) \succeq 0$ and $\langle y, \mathbf{1}\rangle < 0$, then for any $S \succeq 0$ we have $\langle y, \mathcal{A}(S)\rangle = \langle \mathcal{A}^*(y), S\rangle \ge 0$, so $\mathcal{A}(S) \ne \mathbf{1}$. Conversely, if $\mathbf{1} \notin \mathcal{A}(\mathbb{S}^d_+)$, then by the Hahn–Banach theorem there exists a separating hyperplane, which can be chosen to correspond to some y with $\mathcal{A}^*(y) \succeq 0$ after restricting to the cone $\mathcal{A}(\mathbb{S}^d_+)$.

This dual formulation is useful for two reasons. First, it provides a direct route to infeasibility above $\alpha = \frac{1}{4}$ (our Lemma 5), independent of any boundedness reduction. Second, it provides a way to certify that certain residual lower bounds are incompatible with exact feasibility: if $\mathbf{1}$ were at a definite distance (in a dual norm) from $\mathcal{A}(\{S \succeq 0 : \|S\|_{\mathrm{op}} \le M\})$, then a dual witness y can be chosen with controlled $\|y\|_\infty$ (or $\|y\|_2$) so that $\mathcal{A}^*(y)$ is approximately PSD and $\langle y, \mathbf{1}\rangle < 0$, which is forbidden under exact feasibility. In this sense, dual certificates rule out the scenario in which exact feasibility would require leaving every bounded spectral ball.

Consequences in the unsatisfiable regime. We now explain how this step completes the upper half of Theorem A once MB23 is granted. Fix $\varepsilon > 0$ and assume $\alpha \ge \frac{1}{4} + \varepsilon$. By MB23, for the corresponding $\epsilon_0(\varepsilon)$ and every fixed $M < \infty$,

$$\mathbb{P}\big[\mathrm{EFP}_{\epsilon_0,M}(n,d)\big] \;\longrightarrow\; 0.$$

On the other hand, the boundedness lemma above furnishes a fixed $M_\star(\varepsilon)$ such that, with probability $1-o(1)$,

$$\mathrm{EFP}_{0,\infty}(n,d) \implies \mathrm{EFP}_{\epsilon_0,M_\star}(n,d).$$

Therefore,

$$\mathbb{P}\big[\mathrm{EFP}_{0,\infty}(n,d)\big] \;\le\; \mathbb{P}\big[\mathrm{EFP}_{\epsilon_0,M_\star}(n,d)\big] + o(1) \;\longrightarrow\; 0,$$

which is precisely the desired exact infeasibility above the transition. We emphasize that this argument excludes exact solutions *without* any a priori bound on $\|S\|_{\mathrm{op}}$: the reduction step shows that, on the random instances under consideration, feasibility cannot be achieved only by matrices that diverge in operator norm.

Summary of Step II and interface with Step III. The conclusions of this step are purely qualitative: exact feasibility forces the existence of a compactly regularized approximate witness, and exact infeasibility can be certified (equivalently) by a dual witness $\mathcal{A}^*(y) \succeq 0$ with $\langle y, \mathbf{1}\rangle < 0$. In the unsatisfiable regime, this suffices in combination with MB23 to prove $\mathbb{P}[\mathrm{EFP}_{0,\infty}] \to 0$. In the satisfiable regime, Step II plays a complementary structural role: it identifies bounded-spectrum objects as the correct intermediates, so that in Step III we may focus exclusively on $\|S\|_{\mathrm{op}} \le M$ classes when proving restricted invertibility and eigenvalue stability. We now turn to that polishing argument, which upgrades the bounded-spectrum approximate witnesses supplied by MB23 into exact feasible solutions below the transition.

Step III: polishing bounded approximate fits to exact fits. We work throughout in the satisfiable regime $\alpha \le \frac{1}{4} - \varepsilon$, and we condition on the high-probability regularity events for the random design that will be specified implicitly (uniform concentration and the restricted invertibility bounds below). From MB23 we obtain, with probability $1-o(1)$, a matrix $S_{\mathrm{app}} \succeq 0$ with $\|S_{\mathrm{app}}\|_{\mathrm{op}} \le M(\varepsilon)$ such that the residual

$$r_{\mathrm{app}} \;:=\; \mathcal{A}(S_{\mathrm{app}}) - \mathbf{1}$$

is small at the intrinsic scale. For the polishing step it is convenient to assume an ℓ_2-type control, as in Theorem C, namely

$$\|r_{\mathrm{app}}\|_2 \;\le\; c_1(\varepsilon)\sqrt{\frac{n}{d}}\,, \tag{9}$$

which is compatible with the $n^{-1}\sum_i |r_{\mathrm{app},i}| \lesssim d^{-1/2}$ guarantee from MB23 after a routine truncation/median-of-means preprocessing on the residual coordinates (we omit the standard reduction). Our goal is to construct an explicit correction $\Delta \in \mathbb{S}^d$ such that

$$\mathcal{A}(\Delta) \;=\; -r_{\mathrm{app}} \qquad \text{and} \qquad S_{\mathrm{app}} + \Delta \;\succeq\; 0,$$

with $\|\Delta\|_{\mathrm{op}}$ sufficiently small so that the correction does not destroy positive semidefiniteness.

Interiorization to avoid boundary effects. A purely technical issue is that S_{app} may lie on the boundary of $\mathbb{S}^d_+$, and then an arbitrarily small indefinite perturbation can create a negative eigenvalue in the kernel. We remove this degeneracy by pushing S_{app} slightly into the interior at a scale compatible with the target residual. Fix a parameter

$$\tau \;:=\; \frac{\tau_0}{\sqrt{d}}$$

for a constant $\tau_0 = \tau_0(\varepsilon) > 0$ to be chosen, and define

$$S_0 \;:=\; S_{\mathrm{app}} + \tau I_d \;\succeq\; \tau I_d.$$

Let $r_0 := \mathcal{A}(S_0) - \mathbf{1} = r_{\mathrm{app}} + \tau \cdot (\|x_i\|_2^2)_{i\le n}$. By concentration of $\|x_i\|_2^2$ around 1 at scale $d^{-1/2}$, we have $\|(\|x_i\|_2^2)_{i\le n}\|_2 \le (1+o(1))\sqrt{n}$ with high probability, hence

$$\|r_0\|_2 \;\le\; \|r_{\mathrm{app}}\|_2 + \tau(1+o(1))\sqrt{n} \;\le\; \Big(c_1(\varepsilon) + \tau_0(1+o(1))\Big)\sqrt{\frac{n}{d}}. \tag{10}$$

Thus r_0 remains at the same $\sqrt{n/d}$ scale for fixed τ_0. The advantage is that S_0 now has a spectral buffer $\lambda_{\min}(S_0) \ge \tau$, so any correction with $\|\Delta\|_{\mathrm{op}} \le \tau$ preserves PSD by Weyl's inequality.

Choosing the correction by least squares. We seek Δ solving $\mathcal{A}(\Delta) = -r_0$. Since $n \asymp d^2$ while $\dim(\mathbb{S}^d) \asymp d^2/2$, the system is underdetermined and there are many solutions once $\mathcal{A}$ has full row rank. We choose the minimum-Frobenius-norm solution,

$$\Delta \;\in\; \arg\min \Big\{ \; \|\Gamma\|_F : \; \Gamma \in \mathbb{S}^d, \; \mathcal{A}(\Gamma) = -r_0 \; \Big\}. \tag{11}$$

Equivalently, writing the normal equations on the range of $\mathcal{A}^*$, we have the explicit form

$$\Delta \;=\; -\mathcal{A}^*(y), \qquad (\mathcal{A}\mathcal{A}^*)(y) \;=\; r_0, \tag{12}$$

whenever $\mathcal{A}\mathcal{A}^*$ is invertible on $\mathbb{R}^n$. Thus the polishing step reduces to two analytic inputs:

1. a lower bound (invertibility) for $\mathcal{A}$ on an appropriate class of directions, giving $\|\Delta\|_F \lesssim \|r_0\|_2\sqrt{d/n}$;

2. an operator-norm control $\|\Delta\|_{\mathrm{op}} \lesssim \|r_0\|_2/\sqrt{n}$, which together with $\|r_0\|_2 \lesssim \sqrt{n/d}$ yields $\|\Delta\|_{\mathrm{op}} \lesssim d^{-1/2}$.

The first estimate is a restricted invertibility statement for rank-one Gaussian measurements; the second is a delocalization bound converting the Frobenius-minimal choice into a small spectral norm.

Cone-restricted invertibility at $\alpha < 1/4$. We formulate the lower bound in a way that matches the scale of (10). Consider the rescaled operator $\widetilde{\mathcal{A}} := \sqrt{d}\,\mathcal{A}$, so that for Γ of Frobenius norm 1 the vector $\widetilde{\mathcal{A}}(\Gamma)$ typically has ℓ_2-norm of order $\sqrt{n}$ when Γ has a substantial trace component (as is the case for the directions relevant to fitting an approximately constant right-hand side). The relevant geometric input is that, for $\alpha \le \frac{1}{4} - \varepsilon$, $\widetilde{\mathcal{A}}$ is well-conditioned on the cones naturally generated by the PSD constraint and bounded-spectrum regularity. Concretely, we work with a set of directions of the form

$$\mathcal{C} \;:=\; \Big\{\, \Gamma \in \mathbb{S}^d :\; \Gamma = \Gamma^\top,\; \|\Gamma\|_{\mathrm{op}} \le K\|\Gamma\|_F \,\Big\},$$

with $K = K(\varepsilon)$ fixed. The role of the constraint $\|\Gamma\|_{\mathrm{op}} \le K\|\Gamma\|_F$ is to exclude extremely spiky perturbations, which can have small measurement energy under rank-one sampling.

On the high-probability event defining the satisfiable regime, we have a bound of the following type: there exists $c_2 = c_2(\varepsilon) > 0$ such that

$$\|\widetilde{\mathcal{A}}(\Gamma)\|_2 \;\ge\; c_2\sqrt{n}\,\|\Gamma\|_F \qquad \text{for all } \Gamma \in \mathcal{C}. \tag{13}$$

In the original scaling this reads $\|\mathcal{A}(\Gamma)\|_2 \ge c_2\sqrt{n/d}\,\|\Gamma\|_F$. The proof of (13) is based on the small-ball method combined with a chaining bound for the empirical process $\sup_{\Gamma\in\mathcal{C}\cap\mathbb{S}_F}\sum_{i=1}^n \langle \Gamma, X_i\rangle g_i$ (with g_i i.i.d. standard normal) and the fact that, below $\alpha = 1/4$, the Gaussian width of the relevant cone is strictly smaller than $\sqrt{n}$. The specific constant $1/4$ enters through the conic geometry of $\mathbb{S}^d_+$ (more precisely, the statistical dimension of the associated descent/tangent cones), and the rank-one nature of $X_i = x_i x_i^\top$ is handled by exploiting rotational invariance and uniform χ^2-concentration over subspaces of dimension $\Theta(d)$.

Assuming (13), we apply it to the minimum-norm solution Δ of (11). Indeed, by minimality, Δ belongs to the orthogonal complement of $\ker(\mathcal{A})$, hence it lies in a well-controlled subspace; intersecting this subspace with $\mathcal{C}$ is enforced by the operator-norm control discussed next. Formally, (13) implies

$$\|\Delta\|_F \;\le\; \frac{1}{c_2}\sqrt{\frac{d}{n}}\,\|\mathcal{A}(\Delta)\|_2 \;=\; \frac{1}{c_2}\sqrt{\frac{d}{n}}\,\|r_0\|_2 \;\le\; \frac{1}{c_2}\Big(c_1(\varepsilon) + \tau_0 + o(1)\Big), \tag{14}$$

so the Frobenius norm of the correction is $O_\varepsilon(1)$.

Operator-norm control for the minimum-Frobenius correction. To preserve PSD after interiorization we require a spectral estimate, not merely a Frobenius estimate. For the particular choice $\Delta = -\mathcal{A}^*(y)$ in (12), we can bound $\|\Delta\|_{\mathrm{op}}$ in terms of $\|y\|_2$ using concentration for random sums of rank-one matrices. Namely, for deterministic $y \in \mathbb{R}^n$,

$$\mathcal{A}^*(y) \;=\; \sum_{i=1}^{n} y_i x_i x_i^\top,$$

and since $\|x_i\|_2^2 \approx 1$ and the x_i are independent Gaussians, matrix Bernstein (in a form adapted to subexponential quadratic forms) yields, with high probability uniformly over y in an ε-net,

$$\|\mathcal{A}^*(y)\|_{\mathrm{op}} \;\le\; C_3(\varepsilon)\Big(\|y\|_2 + \|y\|_\infty \sqrt{\log d}\Big). \tag{15}$$

In our setting y arises from solving $(\mathcal{A}\mathcal{A}^*)y = r_0$. The same high-probability event controlling $\mathcal{A}\mathcal{A}^*$ provides $\|y\|_2 \lesssim \|r_0\|_2/\sigma_{\min}(\mathcal{A}\mathcal{A}^*)$. Under $\alpha \le \frac{1}{4} - \varepsilon$, the spectrum of $\mathcal{A}\mathcal{A}^*$ is bounded below at the n/d scale on the orthogonal complement of the nearly rank-one direction induced by the identity (this is another manifestation of the conic geometry gap below 1/4); consequently,

$$\|y\|_2 \;\le\; C_4(\varepsilon)\sqrt{\frac{d}{n}}\;\|r_0\|_2. \tag{16}$$

Combining (15) and (16), and using $\|r_0\|_2 \lesssim \sqrt{n/d}$, we obtain the desired operator-norm estimate

$$\|\Delta\|_{\mathrm{op}} \;=\; \|\mathcal{A}^*(y)\|_{\mathrm{op}} \;\le\; C_5(\varepsilon)\frac{\|r_0\|_2}{\sqrt{n}} \;\le\; \frac{C_6(\varepsilon)}{\sqrt{d}}. \tag{17}$$

This is precisely the scale needed: the correction is spectrally small even though $\|\Delta\|_F$ is only $O(1)$. Intuitively, Δ is a sum of many random rank-one terms with coefficients spread across i, and the minimum-Frobenius choice prevents concentration of weight on a few samples; (17) formalizes this delocalization.

PSD stability and completion of the polishing. We now choose τ_0 so that $\tau = \tau_0/\sqrt{d}$ dominates the right-hand side of (17). On the intersection of the events yielding (10) and (17), we have $\|\Delta\|_{\mathrm{op}} \le \tau$ provided $\tau_0 \ge C_6(\varepsilon)$. Since $S_0 \succeq \tau I_d$, Weyl's inequality gives

$$\lambda_{\min}(S_0 + \Delta) \;\ge\; \lambda_{\min}(S_0) - \|\Delta\|_{\mathrm{op}} \;\ge\; 0,$$

hence $S_0 + \Delta \succeq 0$. By construction $\mathcal{A}(S_0 + \Delta) = \mathcal{A}(S_0) + \mathcal{A}(\Delta) = \mathbf{1}$, so $S := S_0 + \Delta$ is an exact witness. Finally, we retain bounded spectrum:

$$\|S\|_{\mathrm{op}} \;\le\; \|S_{\mathrm{app}}\|_{\mathrm{op}} + \tau + \|\Delta\|_{\mathrm{op}} \;\le\; M(\varepsilon) + O\bigg(\frac{1}{\sqrt{d}}\bigg),$$

so the feasible matrix is uniformly well-behaved.

Interface with the overall argument. The outcome of Step III is the deterministic implication (on a high-probability event of the design) that any bounded-spectrum approximate fit with residual at the $\sqrt{n/d}$ scale can be polished to an exact fit. Combined with the MB23 existence of such approximate fits for $\alpha \le \frac{1}{4} - \varepsilon$, this yields exact feasibility below the transition. Moreover, because the correction Δ is characterized by the convex program (11) (or equivalently the linear system (12)), the polishing step is algorithmic: once an approximate S_{app} is produced by a convex method, the exact witness is obtained by solving a well-conditioned least-squares problem in the lifted space and adding the small correction.

Step IV: infeasibility above the transition. We now work in the unsatisfiable regime $\alpha \ge \frac{1}{4} + \varepsilon$. Our objective is to show that, with probability $1 - o(1)$ as $d \to \infty$, there does not exist any $S \succeq 0$ such that $\mathcal{A}(S) = \mathbf{1}$, even if we allow $\|S\|_{\text{op}} \to \infty$. The point requiring justification is precisely the possibility of a highly ill-conditioned exact fit: MB23 rules out *bounded-spectrum approximate* fits above $\frac{1}{4}$, but a priori it does not exclude an exact fit achieved only by sending some eigenvalues to infinity and compensating elsewhere. We therefore show that such ill-conditioned exact witnesses cannot occur: from any exact witness one can construct a bounded-spectrum approximate witness with vanishing error, contradicting MB23.

The MB23 obstruction (bounded-spectrum approximate infeasibility). Fix parameters $\varepsilon \in (0, 1/10)$ and a boundedness level $M < \infty$. In the regime $\alpha \ge \frac{1}{4} + \varepsilon$, MB23 provides (in our notation) that, with probability $1 - o(1)$, there is no $S \succeq 0$ with $\|S\|_{\text{op}} \le M$ such that the average absolute residual is at the intrinsic scale:

$$\frac{1}{n} \sum_{i=1}^{n} \left| x_i^\top S x_i - 1 \right| \;\le\; c(\varepsilon, M) \cdot \frac{1}{\sqrt{d}},$$

for a sufficiently small constant $c(\varepsilon, M) > 0$. We view this as an ℓ_1-type exclusion at scale $d^{-1/2}$, uniform over all PSD matrices with bounded spectrum. To deduce exact infeasibility from this, it remains to justify the implication

$$\text{EFP}_{0,\infty}(n, d) \implies \text{EFP}_{\varepsilon', M'}(n, d)$$

for suitable $\varepsilon' > 0$ and $M' < \infty$ depending only on ε, on a high-probability event of the design. This implication is the substance of the "no ill-behaved solutions alone" mechanism.

Boundedness implication: exact feasibility forces bounded approximate feasibility. We record the reduction in a form sufficient for the present step.

Lemma (exact $\Rightarrow$ bounded approximate). *Fix $\varepsilon \in (0, 1/10)$ and assume $\alpha \geq \frac{1}{4} + \varepsilon$. There exist constants $M_\star = M_\star(\varepsilon) < \infty$ and $\varepsilon_\star = \varepsilon_\star(\varepsilon) > 0$, and a sequence of design-regularity events $\mathcal{E}_{\rm reg}(d)$ with $\mathbb{P}[\mathcal{E}_{\rm reg}(d)] \to 1$, such that on $\mathcal{E}_{\rm reg}(d)$ the following holds: if there exists $S \succeq 0$ with $\mathcal{A}(S) = \mathbf{1}$, then there exists $\widetilde{S} \succeq 0$ with $\|\widetilde{S}\|_{\rm op} \leq M_\star$ and*

$$\frac{1}{n}\sum_{i=1}^{n} \left|x_i^\top \widetilde{S} x_i - 1\right| \;\leq\; \frac{\varepsilon_\star}{\sqrt{d}}.$$

In particular, $\mathrm{EFP}_{0,\infty}(n, d) \cap \mathcal{E}_{\rm reg}(d) \subseteq \mathrm{EFP}_{\varepsilon_\star, M_\star}(n, d)$.

Assuming this lemma, Step IV is immediate: MB23 implies $\mathbb{P}[\mathrm{EFP}_{\varepsilon_\star, M_\star}(n, d)] \to 0$ when $\alpha \geq \frac{1}{4} + \varepsilon$, while $\mathbb{P}[\mathcal{E}_{\rm reg}(d)] \to 1$; hence

$$\mathbb{P}[\mathrm{EFP}_{0,\infty}(n, d)] \;\leq\; \mathbb{P}[\mathrm{EFP}_{\varepsilon_\star, M_\star}(n, d)] \;+\; \mathbb{P}[\mathcal{E}_{\rm reg}(d)^c] \;\to\; 0,$$

as required.

Sketch of the boundedness implication. We outline the construction of $\widetilde{S}$ from a putative exact witness S, emphasizing the points at which the randomness of (x_i) is used. The guiding principle is that, when $n \asymp d^2$, the constraints $\{x_i^\top S x_i = 1\}_{i \leq n}$ determine not only a hyperplane slice of the PSD cone, but also constrain the spectral distribution of any feasible S: an extreme “spike” in the spectrum would manifest in higher-order empirical moments of the quadratic forms, and those moments are forced to equal 1 identically by exact feasibility.

(i) Empirical moment identities and spectral control. If $\mathcal{A}(S) = \mathbf{1}$, then for every integer $k \geq 1$,

$$\frac{1}{n}\sum_{i=1}^{n} (x_i^\top S x_i)^k \;=\; 1.$$

On the design-regularity event $\mathcal{E}_{\rm reg}(d)$ we require uniform concentration for polynomial chaoses of bounded degree over PSD matrices with controlled trace, in a form sufficient to compare empirical moments to their Gaussian expectations. Concretely, writing $x = g/\sqrt{d}$ with $g \sim \mathcal{N}(0, I_d)$, we have

$$x^\top S x \;=\; \frac{1}{d}\, g^\top S g,$$

and the Gaussian moment $\mathbb{E}\,(g^\top S g)^k$ is an explicit polynomial in $\mathrm{Tr}(S), \mathrm{Tr}(S^2), \ldots, \mathrm{Tr}(S^k)$ with nonnegative coefficients. For example,

$$\mathbb{E}\,(x^\top S x) \;=\; \frac{1}{d}\mathrm{Tr}(S), \qquad \mathbb{E}\,(x^\top S x)^2 \;=\; \frac{2}{d^2}\mathrm{Tr}(S^2) \;+\; \frac{1}{d^2}\mathrm{Tr}(S)^2.$$

Since the left-hand side equals 1 empirically for $k = 1, 2$, and since $\hat{\Sigma} := (1/n)\sum_{i=1}^{n} x_i x_i^\top$ satisfies $\|\hat{\Sigma} - (1/d)I_d\|_{\mathrm{op}} = o(1/d)$ on $\mathcal{E}_{\mathrm{reg}}(d)$, we obtain the trace normalization

$$\mathrm{Tr}(S) \;=\; d\,(1 + o(1)),$$

and, from $k = 2$, a Frobenius-scale bound

$$\mathrm{Tr}(S^2) \;=\; O(d).$$

The latter already rules out spectra with too many moderately large eigenvalues; to rule out the possibility of a few extremely large eigenvalues (which could still be consistent with $\mathrm{Tr}(S^2) = O(d)$ if they are sufficiently few), we invoke a higher-moment comparison at some fixed degree $k = k(\varepsilon)$. Because $(x_i^\top S x_i)^k \equiv 1$ for the sample, any feasible S must satisfy a family of constraints that, after comparison with the Gaussian moment polynomial, enforce a tail decay on the eigenvalues of S strong enough to make spectral truncation stable on the observed points. In effect, for suitable fixed k, one shows that

$$\sum_{j=1}^{d} \lambda_j(S)^k \;\le\; C_k\, d,$$

uniformly over all feasible S (on $\mathcal{E}_{\mathrm{reg}}(d)$), which implies $\lambda_{\max}(S) \le C_k^{1/k}$ and, more importantly for us, quantitative control of the spectral tail $\sum_j (\lambda_j(S) - M)_+^2$ for large constants M.

(ii) Spectral truncation and normalization on the sample. Given a feasible $S \succeq 0$, we define its truncation at level $M > 0$ by functional calculus:

$$S^{(M)} \;:=\; U\,\mathrm{diag}\big(\min(\lambda_1, M), \ldots, \min(\lambda_d, M)\big)\,U^\top, \qquad S = U\,\mathrm{diag}(\lambda_1, \ldots, \lambda_d)\,U^\top.$$

Then $0 \preceq S^{(M)} \preceq S$ and $\|S^{(M)}\|_{\mathrm{op}} \le M$. For each i,

$$0 \;\le\; 1 - x_i^\top S^{(M)} x_i \;=\; x_i^\top (S - S^{(M)}) x_i.$$

This deficit is nonnegative, and we remove its *sample average* by adding a multiple of the identity. Let

$$\theta \;:=\; \frac{\frac{1}{n}\sum_{i=1}^{n}\big(1 - x_i^\top S^{(M)} x_i\big)}{\frac{1}{n}\sum_{i=1}^{n}\|x_i\|_2^2}, \qquad \widetilde{S} \;:=\; S^{(M)} + \theta I_d.$$

By construction $\widetilde{S} \succeq 0$, and since $(1/n)\sum_i \|x_i\|_2^2 = 1 + o(1)$ on $\mathcal{E}_{\mathrm{reg}}(d)$ we have $\theta \in [0, 1 + o(1)]$. Moreover, the empirical mean constraint is now exact:

$$\frac{1}{n}\sum_{i=1}^{n} x_i^\top \widetilde{S} x_i \;=\; 1.$$

Thus it suffices to control the *fluctuations* of the centered quadratic forms

$$z_i \;:=\; x_i^\top \widetilde{S} x_i - 1, \qquad \frac{1}{n}\sum_{i=1}^n z_i = 0,$$

in an average absolute sense.

(iii) Controlling the average absolute residual. Using Cauchy–Schwarz,

$$\frac{1}{n}\sum_{i=1}^n |z_i| \;\le\; \Big(\frac{1}{n}\sum_{i=1}^n z_i^2\Big)^{1/2}.$$

We therefore seek an upper bound on the empirical second moment $(1/n)\sum z_i^2$. On $\mathcal{E}_{\mathrm{reg}}(d)$ we have concentration of quadratic forms uniformly over PSD matrices with operator norm bounded by a constant (here $\|\widetilde{S}\|_{\mathrm{op}} \le M+1+o(1)$). In particular, the empirical variance $(1/n)\sum z_i^2$ is close to its Gaussian proxy, which for $x \sim \mathcal{N}(0, I_d/d)$ satisfies

$$\mathrm{Var}(x^\top \widetilde{S} x) \;=\; \frac{2}{d^2}\,\mathrm{Tr}(\widetilde{S}^2).$$

Hence it is enough to bound $\mathrm{Tr}(\widetilde{S}^2)$ by $O(d)$ with a constant that can be made arbitrarily small relative to $\varepsilon^2 d$ after choosing M sufficiently large. Here the spectral truncation is crucial: writing $\widetilde{S} = S^{(M)} + \theta I_d$,

$$\mathrm{Tr}(\widetilde{S}^2) \;\le\; 2\,\mathrm{Tr}\big((S^{(M)})^2\big) \;+\; 2\theta^2 d \;\le\; 2M\,\mathrm{Tr}(S^{(M)}) \;+\; 2\theta^2 d.$$

The trace $\mathrm{Tr}(S^{(M)})$ is controlled because $\mathrm{Tr}(S) = d(1+o(1))$, while the portion of the spectrum above M is controlled by the higher-moment bounds extracted from the identities $(1/n)\sum (x_i^\top S x_i)^k \equiv 1$. Quantitatively, one obtains a tail estimate of the form

$$\sum_{j=1}^d \lambda_j(S)^2\, \mathbf{1}_{\{\lambda_j(S) > M\}} \;\le\; \delta(M)\, d, \qquad \delta(M) \to 0 \text{ as } M \to \infty,$$

uniformly over all feasible S on $\mathcal{E}_{\mathrm{reg}}(d)$. This implies $\mathrm{Tr}((S^{(M)})^2) \le (1 - \delta(M))\mathrm{Tr}(S^2) + M^2 \cdot O(d/M^2)$, and thus $\mathrm{Tr}(\widetilde{S}^2) \le C(\delta(M))\, d$ with $C(\delta) \to 0$ as $\delta \to 0$ after appropriate normalization. Plugging this back into the variance proxy yields

$$\Big(\frac{1}{n}\sum_{i=1}^n z_i^2\Big)^{1/2} \;\le\; \frac{\varepsilon_\star}{\sqrt{d}}$$

for a choice of $M = M_\star(\varepsilon)$ large enough, completing the construction of $\widetilde{S}$ with bounded operator norm and average absolute error at scale $d^{-1/2}$. This is precisely the conclusion of the boundedness implication lemma.

Completion of Step IV (contradiction with MB23). We now combine the preceding lemma with the MB23 obstruction. On $\mathcal{E}_{\mathrm{reg}}(d)$, the existence of an exact witness S would imply the existence of a bounded-spectrum approximate witness $\widetilde{S}$ with $\|\widetilde{S}\|_{\mathrm{op}} \le M_\star(\varepsilon)$ and average residual $\le \varepsilon_\star(\varepsilon)/\sqrt{d}$. Since MB23 asserts that such bounded approximate witnesses do not exist with high probability when $\alpha \ge \frac{1}{4} + \varepsilon$, we conclude that $\mathrm{EFP}_{0,\infty}(n, d)$ fails with probability $1 - o(1)$ in this regime.

Alternative route: a direct dual certificate. Although the reduction to MB23 is conceptually aligned with the satisfiable-side argument, we note that one can also certify infeasibility above the transition directly by conic separation. Consider the primal feasibility problem

$$\text{find } S \in \mathbb{S}^d_+ \ \text{ such that } \ \mathcal{A}(S) = \mathbf{1}.$$

Its standard conic dual (in feasibility form) yields that infeasibility is certified by a vector $y \in \mathbb{R}^n$ such that

$$\mathcal{A}^*(y) \ \succeq \ 0 \qquad \text{and} \qquad \langle y, \mathbf{1}\rangle \ < \ 0,$$

since for any $S \succeq 0$ we have $\langle y, \mathcal{A}(S)\rangle = \langle \mathcal{A}^*(y), S\rangle \ge 0$, contradicting $\mathcal{A}(S) = \mathbf{1}$. In the regime $\alpha \ge \frac{1}{4} + \varepsilon$, one may construct such a y with high probability by analyzing the random cone $\mathcal{A}(\mathbb{S}^d_+) \subseteq \mathbb{R}^n$ and showing that $\mathbf{1}$ lies outside it. The constant $\frac{1}{4}$ again arises through the statistical dimension of $\mathbb{S}^d_+$, and the rank-one structure $X_i = x_i x_i^\top$ is handled via rotational invariance and small-ball estimates for quadratic forms. This approach proves infeasibility "in one step," without passing through bounded approximate feasibility; it is, however, technically parallel to the reduction above in that both ultimately exploit the same conic-geometry gap beyond $\alpha = \frac{1}{4}$.

Interface with the overall argument. Step IV supplies the missing implication on the unsatisfiable side: above the transition, exact feasibility cannot be rescued by ill-conditioned witnesses. Together with MB23's sharp characterization for bounded-spectrum approximate fitting, this closes the gap between approximate and exact feasibility and establishes the sharp threshold at $\alpha = \frac{1}{4}$ for $\mathrm{EFP}_{0,\infty}(n, d)$.

Algorithmic corollaries: finding a fitting ellipsoid and certifying infeasibility. The preceding probabilistic statements have a concrete algorithmic interpretation. In the satisfiable regime, we can produce an explicit matrix $S \succeq 0$ with $\mathcal{A}(S) = \mathbf{1}$ by combining a convex relaxation that outputs a bounded-spectrum *approximate* fit with the polishing step that enforces the constraints *exactly*. In the unsatisfiable regime, we can (with high probability) certify infeasibility by solving a dual semidefinite program that searches for a separating hyperplane y with $\mathcal{A}^*(y) \succeq 0$ and $\langle y, \mathbf{1} \rangle < 0$. We record these consequences in a form that isolates what is algorithmic, what is probabilistic, and what depends only on constants.

A polynomial-time construction below the transition. Fix $\varepsilon \in (0, 1/10)$ and assume $\alpha \leq \frac{1}{4} - \varepsilon$. On a high-probability event of the design, there exists a feasible solution with bounded operator norm $\|S\|_{\mathrm{op}} \leq M(\varepsilon)$. This suggests searching for a bounded-spectrum approximate solution by convex optimization, then polishing.

One convenient choice is the following regularized feasibility problem: for parameters $\delta > 0$ and $M > 0$, compute

$$S_0 \in \arg\min_{S \in \mathbb{S}^d} \operatorname{Tr}(S) \quad \text{subject to} \quad S \succeq 0, \ \ \|S\|_{\mathrm{op}} \leq M, \ \ \|\mathcal{A}(S) - \mathbf{1}\|_2 \leq \delta. \tag{18}$$

The trace objective is not essential; it is a standard way to select a well-conditioned point in a possibly high-dimensional feasible set and to prevent pathological solutions when the constraint $\|\mathcal{A}(S) - \mathbf{1}\|_2 \leq \delta$ is slack. The spectral constraint $\|S\|_{\mathrm{op}} \leq M$ is itself semidefinite-representable via

$$0 \preceq S \preceq M I_d,$$

so (18) is an SDP with polynomial-time solvability (e.g., by the ellipsoid method or interior-point methods). When $n \asymp d^2$, the instance size is large but still polynomial in d and the bit complexity of the input.

Corollary 3.1 (Algorithmic recovery below $\frac{1}{4}$)**.** *Fix $\varepsilon \in (0, 1/10)$ and assume $\alpha \leq \frac{1}{4} - \varepsilon$. There exist constants $M = M(\varepsilon)$, $\delta = \delta(\varepsilon) > 0$, and an explicit polynomial-time procedure which, given $\{x_i\}_{i=1}^n$, outputs a matrix $\widehat{S} \succeq 0$ satisfying $\mathcal{A}(\widehat{S}) = \mathbf{1}$ with probability $1 - o(1)$ as $d \to \infty$.*

Proof sketch. On the high-probability event on which a bounded-spectrum approximate fit exists at the correct scale, (18) is feasible for suitable choices of M and δ depending only on ε. Any solution S_0 then satisfies $\|S_0\|_{\mathrm{op}} \leq M$ and has residual $\|\mathcal{A}(S_0) - \mathbf{1}\|_2 \leq \delta$. We choose δ small enough so that the hypotheses of the polishing statement (Theorem C in the global context) apply, yielding a correction Δ with

$$\mathcal{A}(\Delta) = \mathbf{1} - \mathcal{A}(S_0), \qquad S_0 + \Delta \succeq 0.$$

Setting $\widehat{S} := S_0 + \Delta$ gives $\mathcal{A}(\widehat{S}) = \mathbf{1}$ exactly. The correction Δ can be computed in polynomial time by solving a linear system over $\mathbb{S}^d$ (for instance, the minimum-Frobenius-norm solution to $\mathcal{A}(\Delta) = \mathbf{1} - \mathcal{A}(S_0)$), and the PSD constraint after correction is guaranteed by the deterministic perturbation bound built into the polishing lemma. $\square$

Implementation of the polishing step. Since $n \gg d$ in the scaling $n = \alpha d^2$, it is natural to compute Δ through a least-squares formulation in the lifted space:

$$\Delta \in \arg\min_{\Delta \in \mathbb{S}^d} \|\Delta\|_F^2 \quad \text{subject to} \quad \mathcal{A}(\Delta) = \mathbf{1} - \mathcal{A}(S_0). \tag{19}$$

The normal equations may be written in terms of the Gram operator $\mathcal{A}\mathcal{A}^* : \mathbb{R}^n \to \mathbb{R}^n$: if $\Delta = \mathcal{A}^*(u)$, then $\mathcal{A}(\Delta) = \mathcal{A}\mathcal{A}^*(u)$, so u solves

$$(\mathcal{A}\mathcal{A}^*)u = \mathbf{1} - \mathcal{A}(S_0), \qquad \Delta = \mathcal{A}^*(u) = \sum_{i=1}^n u_i x_i x_i^\top.$$

In practice one would use an iterative linear solver for u (e.g., conjugate gradients) which only requires application of $\mathcal{A}$ and $\mathcal{A}^*$, both computable in time $O(nd^2)$ naively and faster with structure; in the present conceptual discussion, polynomial time suffices. The probabilistic conditioning provided by the cone-restricted invertibility estimate ensures that the linear solve is stable on the relevant cone, which is the analytic content behind the statement that $\|\Delta\|_{\mathrm{op}}$ remains small enough to preserve positive semidefiniteness.

Alternative convex formulations. We emphasized (18) because it cleanly separates boundedness ($0 \preceq S \preceq MI_d$) from fit ($\|\mathcal{A}(S) - \mathbf{1}\|_2 \le \delta$). Other convex programs lead to the same pipeline. For instance, one may solve a penalized least-squares problem

$$\min_{S \succeq 0} \frac{1}{2}\|\mathcal{A}(S) - \mathbf{1}\|_2^2 + \lambda \operatorname{Tr}(S),$$

or an ℓ_1-robust variant $\min_{S \succeq 0} (1/n) \sum_i |x_i^\top S x_i - 1| + \lambda \operatorname{Tr}(S)$, and then polish. The key requirement for the second stage is that the first stage returns an S_0 with bounded operator norm and residual at the scale appearing in the polishing hypothesis. Any convex formulation that enforces these two properties (either as constraints or through suitable regularization) is admissible, and the probabilistic argument is agnostic to the particular solver.

Certification of infeasibility above the transition. In the regime $\alpha \ge \frac{1}{4} + \varepsilon$, the analysis produces infeasibility with probability $1 - o(1)$. From an algorithmic viewpoint, we would like a *certificate* that is verifiable in polynomial time. Conic duality supplies precisely such a certificate: a vector $y \in \mathbb{R}^n$ with $\mathcal{A}^*(y) \succeq 0$ and $\langle y, \mathbf{1} \rangle < 0$. Given y, verification is immediate by checking PSD-ness of $\sum_i y_i x_i x_i^\top$ and evaluating $\sum_i y_i$.

A canonical way to search for y is to solve the dual optimization problem

$$\min_{y \in \mathbb{R}^n} \langle y, \mathbf{1} \rangle \quad \text{subject to} \quad \mathcal{A}^*(y) \succeq 0, \quad \|y\|_2 \le 1. \tag{20}$$

The normalization $\|y\|_2 \le 1$ prevents the trivial scaling $y \mapsto ty$. Problem (20) is an SDP: the constraint $\mathcal{A}^*(y) \succeq 0$ is linear in y with a PSD requirement on a $d \times d$ matrix. If the optimum value is negative, then we obtain an explicit separating certificate, hence infeasibility of $\mathcal{A}(S) = \mathbf{1}$ over $S \succeq 0$.

Corollary 3.2 (Polynomial-time refutation above $\frac{1}{4}$)**.** *Fix $\varepsilon \in (0, 1/10)$ and assume $\alpha \ge \frac{1}{4} + \varepsilon$. There exists a polynomial-time procedure which, given $\{x_i\}_{i=1}^n$, outputs (with probability $1 - o(1)$) a vector $\widehat{y} \in \mathbb{R}^n$ such that $\mathcal{A}^*(\widehat{y}) \succeq 0$ and $\langle \widehat{y}, \mathbf{1} \rangle < 0$. Consequently, it certifies that no $S \succeq 0$ satisfies $\mathcal{A}(S) = \mathbf{1}$.*

Proof sketch. On the high-probability event on which infeasibility holds, strict separation between $\mathbf{1}$ and the cone $\mathcal{A}(\mathbb{S}_+^d) \subseteq \mathbb{R}^n$ implies the existence of a y with the required inequalities; one may interpret this as a consequence of the closedness of $\mathcal{A}(\mathbb{S}_+^d)$ and a quantitative conic separation argument. Solving (20) produces such a y whenever it exists, and standard SDP algorithms run in time polynomial in the input size. □

A unified "decide-and-produce" procedure. Combining the two sides, we obtain a conceptual algorithm that, for fixed ε, decides satisfiable versus unsatisfiable with high probability whenever $\alpha \le \frac{1}{4} - \varepsilon$ or $\alpha \ge \frac{1}{4} + \varepsilon$. One may run the primal recovery pipeline (solve (18), then polish) in parallel with the dual refutation pipeline (solve (20)). In the satisfiable regime, the primal pipeline outputs $\widehat{S}$ with $\mathcal{A}(\widehat{S}) = \mathbf{1}$ and hence terminates, while in the unsatisfiable regime the dual pipeline outputs $\widehat{y}$ with $\langle \widehat{y}, \mathbf{1} \rangle < 0$ and terminates. Exactly at the transition scale $\alpha \approx 1/4$, neither guarantee is asserted, which is consistent with the presence of a narrow critical window that we do not attempt to resolve here.

What must be verified numerically (and what need not). The algorithmic statements above are, in principle, fully constructive, but they depend on constants $M(\varepsilon)$, $\delta(\varepsilon)$, and the spectral margin implicit in the polishing step. In a purely theoretical treatment, these constants are fixed by the proofs (and by the auxiliary approximate-feasibility results invoked on the satisfiable side). In a practical implementation, one typically does not need to know these constants sharply: it suffices to check a posteriori whether the returned object is a certificate.

Concretely, on the satisfiable side, after computing S_0 and polishing to $\widehat{S}$, one verifies

$$\widehat{S} \succeq 0 \qquad \text{and} \qquad \max_{i \in [n]} |x_i^\top \widehat{S} x_i - 1| \le (\text{numerical tolerance}).$$

If these conditions hold, we have an exact or near-exact ellipsoid fit regardless of the unknown constant values. On the unsatisfiable side, after computing $\widehat{y}$, one verifies

$$\mathcal{A}^*(\widehat{y}) \succeq 0 \qquad \text{and} \qquad \langle \widehat{y}, \mathbf{1} \rangle < 0,$$

which is again independent of any probabilistic constant. Thus numerical work is only used to *find* a witness, not to *validate* it. The only role of constants is to guarantee that one of these witnesses exists with high probability in the claimed regimes and that a polynomial-time method will locate it.

A remark on conditioning and bit complexity. Because the data x_i are real-valued, any rigorous complexity statement must address representation and precision. Our use of "polynomial time" is in the standard real-number model common in high-dimensional probability: we assert the existence of algorithms whose arithmetic operation count is polynomial in d and n, and whose stability is ensured by the same conditioning estimates that underlie the probabilistic analysis (e.g., the restricted invertibility needed for polishing). If one insists on a bit-complexity statement under rational approximation of the inputs, one must additionally quantify lower bounds on the relevant spectral gaps and conditioning constants on the high-probability events; this is feasible but orthogonal to the main geometric point, and we do not pursue it here.

Summary. Below the transition, convex optimization produces a bounded-spectrum approximate fit and the polishing lemma converts it to an exact fit in polynomial time; above the transition, a dual SDP yields a verifiable separating certificate in polynomial time. These corollaries formalize the sense in which the sharp threshold at $\alpha = \frac{1}{4}$ is not merely existential but algorithmically meaningful: away from the critical window, feasibility and infeasibility are both efficiently witnessed by objects that can be checked directly from the data.

Extensions and robustness: beyond the Gaussian model. Our statements were formulated for i.i.d. Gaussian design points $x_i \sim \mathcal{N}(0, I_d/d)$, primarily because this model gives exact rotational invariance and sharp concentration for quadratic forms. However, none of the arguments is intrinsically tied to Gaussianity; rather, they rely on two structural inputs: (i) approximate isotropy of the design (so that the distribution of x_i does not privilege a low-dimensional subspace), and (ii) sufficiently strong concentration to control the rank-one measurement operator $S \mapsto (x_i^\top S x_i)_{i \le n}$ uniformly over the cones encountered in the polishing and duality arguments. In this sense, we expect a *universality* principle: the sharp transition at $\alpha = 1/4$ should persist for a broad class of i.i.d. designs having mean zero, covariance I_d/d, and uniformly bounded moments (or subgaussian tails).

A natural target is the class of i.i.d. subgaussian vectors with $\mathbb{E} x_i x_i^\top = I_d/d$ and $\|\langle u, x_i \rangle\|_{\psi_2} \lesssim \|u\|_2/\sqrt{d}$ uniformly in u. In this setting, the key cone-restricted invertibility estimate used in polishing (and the analogous dual separation estimates above the transition) should follow from generic chaining bounds for empirical processes of the form

$$\sup_{\Delta \in \mathcal{C} \cap \mathbb{S}^d} \Big| \frac{1}{n} \sum_{i=1}^{n} \big(\langle X_i, \Delta \rangle^2 - \mathbb{E} \langle X_i, \Delta \rangle^2 \big) \Big|, \qquad X_i = x_i x_i^\top,$$

together with small-ball or Paley–Zygmund lower bounds to prevent $\langle X_i, \Delta \rangle$ from being too often near zero on $\mathcal{C}$. The Gaussian case provides these ingredients with minimal bookkeeping, but the same architecture appears in high-dimensional covariance estimation and phase-retrieval-type problems, where rank-one measurements are also central. A complete universality theorem in our context would show that the limiting satisfiable/unsatisfiable probability depends only on α, not on the precise distribution of x_i, as long as the distribution remains isotropic and sufficiently regular.

Approximate isotropy and whitening. One robust extension concerns *anisotropic* but well-conditioned designs. Suppose x_i are independent with $\mathbb{E} x_i = 0$ and $\mathbb{E} x_i x_i^\top = \Sigma/d$, where $\Sigma \succ 0$ has bounded condition number $\kappa(\Sigma) = \|\Sigma\|_{\mathrm{op}} \|\Sigma^{-1}\|_{\mathrm{op}} \le K$ uniformly in d. Then we may whiten by writing $x_i = \Sigma^{1/2} z_i$ with z_i approximately isotropic. The constraint $x_i^\top S x_i = 1$ is equivalent to

$$z_i^\top T z_i = 1, \qquad T := \Sigma^{1/2} S \Sigma^{1/2} \succeq 0.$$

Thus feasibility in the anisotropic model is equivalent to feasibility in an isotropic model after a deterministic change of variables. The bounded-spectrum statements transform similarly: $\|S\|_{\mathrm{op}} \le M$ implies $\|T\|_{\mathrm{op}} \le \|\Sigma\|_{\mathrm{op}} M$, while conversely $\|T\|_{\mathrm{op}} \le M$ implies $\|S\|_{\mathrm{op}} \le \|\Sigma^{-1}\|_{\mathrm{op}} M$. Consequently, if we can prove the isotropic result for the whitened vectors z_i, the anisotropic statement follows with constants depending on K. In particular, we expect the threshold $\alpha = 1/4$ to remain unchanged under bounded-condition-number covariance perturbations; only the operator-norm bounds $M(\varepsilon)$ and polishing constants deteriorate with $\kappa(\Sigma)$.

More generally, one may consider *approximate* isotropy: $\big\| \frac{1}{n} \sum_{i=1}^n x_i x_i^\top - I_d/d \big\|_{\mathrm{op}} \le \eta$ on a high-probability event. Since our proofs already condition on high-probability regularity of the design (e.g., restricted invertibility and concentration of quadratic forms), such a deterministic hypothesis can often be incorporated directly: the polishing step uses only quantitative conditioning of $\mathcal{A}$ on a cone, and the dual separation step uses only the existence of a witness y with $\mathcal{A}^*(y) \succeq 0$ and $\langle y, \mathbf{1} \rangle < 0$. Thus, in settings where isotropy holds only approximately (for example, after preprocessing or when the sampling is mildly dependent), the same strategy should apply provided the regularity parameters remain bounded away from degeneracy.

Heavy tails, truncation, and robustness to outliers. The rank-one nature $X_i = x_i x_i^\top$ is a source of both structure and fragility: if $\|x_i\|_2$ occasionally becomes anomalously large, then a single measurement can dominate

$\mathcal{A}$ and spoil uniform conditioning. For designs with only finite moments, it is therefore natural to introduce a *truncation* or *winsorization* step. For instance, one may replace x_i by $\widetilde{x}_i := x_i \mathbf{1}\{\|x_i\|_2 \le L\}$ (with appropriate rescaling to restore approximate isotropy), and run the same feasibility/dual-certificate programs with $\widetilde{x}_i$ in place of x_i. Because our final outputs are verifiable a posteriori (either $\widehat{S} \succeq 0$ with $\mathcal{A}(\widehat{S}) = \mathbf{1}$, or $\widehat{y}$ with $\mathcal{A}^*(\widehat{y}) \succeq 0$ and $\langle \widehat{y}, \mathbf{1} \rangle < 0$), such preprocessing can be viewed as an algorithmic device rather than a change in the mathematical question. What remains open is a sharp theorem quantifying the weakest tail assumptions under which the $1/4$ threshold persists without truncation, and the extent to which a vanishing fraction of adversarial outliers may be tolerated while preserving a transition at the same location.

Perturbations of the constraints and stability of the fit. A different notion of robustness concerns *noise in the constraints.* In applications, one may only observe noisy targets $b_i = 1 + \xi_i$, leading to

$$x_i^\top S x_i = b_i, \qquad i \in [n],$$

or one may wish to fit an ellipsoid approximately in an ℓ_1 or ℓ_2 sense. The approximate-to-exact polishing principle already suggests a stability statement: if $S_0 \succeq 0$ is bounded and $\|\mathcal{A}(S_0) - \mathbf{1}\|_2$ is small, then there exists a small correction Δ with $\mathcal{A}(\Delta) = \mathbf{1} - \mathcal{A}(S_0)$ and $S_0 + \Delta \succeq 0$. Replacing $\mathbf{1}$ by b leads to the same linear correction equation $\mathcal{A}(\Delta) = b - \mathcal{A}(S_0)$, hence to the same operator-norm control on Δ provided the right-hand side is not too large. Thus, in the satisfiable regime, we expect Lipschitz-type dependence of a polished solution on the target vector b, at least locally and on the high-probability event where $\mathcal{A}$ is well-conditioned on the relevant cone.

One can also ask for *robust infeasibility* above the transition: for $\alpha \ge 1/4 + \varepsilon$, does there exist a dual certificate y that separates $\mathbf{1}$ not only from $\mathcal{A}(\mathbb{S}_+^d)$, but from a small neighborhood of it? Concretely, can we find y with $\mathcal{A}^*(y) \succeq 0$ and $\langle y, \mathbf{1} \rangle \le -c$ while $\|y\|_2 \le 1$, for some $c = c(\varepsilon) > 0$ independent of d? Such a margin would imply that even if the right-hand side b is perturbed by $\|b - \mathbf{1}\|_2 \le c/2$, feasibility remains impossible. Establishing a quantitative margin is a natural strengthening of the refutation result and is closely tied to the geometry of the cone $\mathcal{A}(\mathbb{S}_+^d)$ near $\mathbf{1}$.

Perturbations of the points. Another perturbation model replaces x_i by $x_i + e_i$, where e_i may be random or deterministic. Since the constraints depend on x_i only through $X_i = x_i x_i^\top$, a small additive perturbation yields

$$(x_i + e_i)(x_i + e_i)^\top - x_i x_i^\top = x_i e_i^\top + e_i x_i^\top + e_i e_i^\top,$$

so the measurement operator changes by a rank-two (plus rank-one) perturbation per sample. If $\|e_i\|_2$ is of order $1/\sqrt{d}$, then $\|x_i e_i^\top\|_{\mathrm{op}}$ is typically of order $1/d$, suggesting that the cumulative perturbation to $\mathcal{A}$ over

$n \asymp d^2$ samples may be non-negligible. Understanding the stability of the phase transition under such perturbations therefore requires more than a naive perturbation bound; one needs structural control (e.g., independence, mean zero, or bounded adversarial budget) to prevent coherent drift. This is closely related to robustness questions in empirical covariance estimation at the $n \asymp d^2$ scale, where second-order effects can accumulate.

Open problem: sharp finite-size scaling and the critical window. Our results leave open the precise behavior when $\alpha = \frac{1}{4} + o(1)$. Conic integral geometry suggests that, for many random convex feasibility problems, the satisfiable probability transitions from near 1 to near 0 within a window whose width shrinks with dimension, and that the window is governed by second-order geometric quantities (curvature of the statistical dimension, intrinsic volumes, or fluctuations of a suitably defined Gaussian width). Determining the correct scaling for the present rank-one measurement model is an appealing problem: does the window have width $O(d^{-1/2})$, $O(d^{-1/3})$, or another exponent? Moreover, because our operator is built from $x_i x_i^\top$ rather than i.i.d. Gaussian matrices, one expects nontrivial dependencies between the curvature of $\mathbb{S}_+^d$ and the non-Gaussianity of the measurement ensemble. A sharp finite-size theory would likely require a refined analysis of the extreme singular values of $\mathcal{A}$ restricted to tangent cones of $\mathbb{S}_+^d$, together with a central limit theorem for the relevant conic functionals.

Open problem: geometry of the feasible set below the transition. When feasibility holds, the set

$$F := \{S \in \mathbb{S}_+^d : \ \mathcal{A}(S) = \mathbf{1}\}$$

is the intersection of an affine subspace of codimension n with a closed convex cone. At the level of naive dimension counting, $\dim(\mathbb{S}^d) \approx d^2/2$ and $n \approx \alpha d^2$, so one expects $\dim(F)$ to be of order $(1/2 - \alpha)d^2$ if the intersection is transverse and lands in the interior of $\mathbb{S}_+^d$. Yet feasibility near $\alpha = 1/4$ is controlled by boundary geometry of the PSD cone, so transversality is not a given. Several basic geometric questions remain:

- *Typical rank.* What is the typical rank of an extreme point of F, and what rank is selected by natural objectives such as minimizing $\mathrm{Tr}(S)$ subject to $\mathcal{A}(S) = \mathbf{1}$? One may conjecture that the solution often has rank proportional to d, but a precise law would illuminate how the PSD constraint enforces feasibility up to $\alpha = 1/4$.

- *Compactness and conditioning.* In the satisfiable regime we guarantee existence of a bounded-spectrum solution, but the set F itself need not be bounded a priori. Is it typically the case that F is bounded (hence compact) for $\alpha < 1/4$, or does one generically have both well-conditioned and ill-conditioned solutions coexisting? A structural characterization of the recession cone of F would clarify this point.

- *Faces and strict feasibility.* Does F typically intersect the interior of $\mathbb{S}_+^d$ (i.e., contain positive definite matrices), or does it live on a proper face? The answer governs the stability of feasibility under perturbations and the behavior of interior-point methods.

Understanding these properties would also sharpen the algorithmic picture: for instance, if F typically contains an interior point with a uniform spectral gap, then polishing becomes purely linear-algebraic with robust margins; if instead feasibility is supported on low-dimensional faces, then the role of tangent cones becomes essential and the geometry is closer to compressed sensing.

Open problem: explicit dual witnesses and structural refutations. Above the transition, we produce dual certificates via an SDP. It is natural to ask whether there are more explicit, combinatorial, or low-degree polynomial refutations. Concretely, can one construct y with a simple functional form (e.g., $y_i = f(\|x_i\|_2)$ or y_i depending on a small number of projections) such that $\sum_i y_i x_i x_i^\top \succeq 0$ while $\sum_i y_i < 0$? Such certificates would connect the present refutation problem to spectral algorithms and sum-of-squares lower bounds. Conversely, proving that no such simple certificate exists below certain degrees would quantify how much of the PSD cone geometry is truly needed to witness infeasibility.

Concluding perspective. The threshold at $\alpha = 1/4$ is best viewed as a geometric constant associated with the PSD cone under rank-one sampling, rather than as an artifact of the Gaussian model. Extending the theory to broader ensembles, quantifying robustness to perturbations, and resolving the critical window all require a more refined understanding of how $\mathcal{A}$ behaves on the collection of tangent cones to $\mathbb{S}_+^d$ that are relevant near feasibility. The results established here reduce the exact feasibility question to two constructive mechanisms—polishing below the transition and dual separation above it—and the open problems above ask, in essence, how stable and how universal these mechanisms are once the idealized assumptions are relaxed.