The Conformal Bootstrap

Henry Purcell

June 2, 2026

Abstract

Conformal field theories (CFTs) arise as fixed points of renormalization group flows and describe universal physics at second-order phase transitions. In $d \geq 3$ dimensions, the conformal bootstrap exploits unitarity and the associativity of the operator product expansion to place bounds on CFT data—namely the operator dimensions and OPE coefficients—nonperturbatively and without reference to any Lagrangian. We review the necessary formalism to understand this idea: the conformal algebra $SO(d+1,1)$ and its representations, conformal correlators, radial quantization, unitarity bounds, the OPE, and conformal blocks. We then derive the crossing symmetry equation for a four-point function of identical scalars, and show how the positivity of OPE coefficients turns it into a constraint on the space of consistent CFTs. We work through the bootstrap bound analytically and present numerical results for identical scalars, reproducing the characteristic kink in the scalar dimension bound that corresponds to the critical 2d Ising model, and discussing generalizations to 3d. We discuss potential future directions for research with the conformal bootstrap.

1 Introduction

One of the central challenges of quantum field theory is extracting physical predictions from strongly coupled systems where perturbation theory fails. A variety of nonperturbative techniques have been developed to achieve this: lattice regularization and strong-coupling expansions reveal confinement in gauge theories, the large N limit exposes planar dominance and simplifies the sum over Feynman diagrams, semiclassical methods such as instantons capture tunneling effects invisible to any order in perturbation theory, and anomalies provide exact constraints on the spectrum and dynamics. Each of these approaches works by exploiting different structural features of the theory that apply regardless of specific dynamics, or at least in some limit.

In this context, conformal field theories (CFTs) arise naturally as fixed points of renormalization group (RG) flows, reached in the deep UV or IR. At such points the theory becomes scale-invariant. The resulting CFT is independent of microscopic details—it depends only on dimensionality and global symmetries—which is the origin of the universality observed at second-order phase transitions. CFTs therefore occupy a distinguished position in the space of quantum field theories: they describe the universal physics at phase transitions and at the endpoints of RG flows.

The conformal symmetry of these theories is extraordinarily constraining. Two- and three-point correlation functions are fixed up to a finite set of numbers: the operator scaling dimensions Δ_i and the OPE coefficients f_{ijk}, collectively known as the *CFT data*. All higher-point correlators can be reduced to these quantities via the operator product expansion (OPE). A CFT is therefore entirely

specified by its CFT data. But this leads to an important question: do all values of this data define consistent theories, and if not, which ones do?

It turns out, in fact, that most choices of CFT data do not produce valid theories. The *conformal bootstrap* is a program for answering this question using only general consistency conditions, without reference to any Lagrangian. The idea, originating in the work of Ferrara, Gatto, and Grillo [1] and Polyakov [2] in the 1970s and revived by Rattazzi, Rychkov, Tonni, and Vichi [3] in 2008, is to exploit *crossing symmetry*—the requirement that the OPE, applied in different channels of a four-point function, must give the same result. Combined with unitarity (which forces OPE coefficients squared to be non-negative), crossing symmetry yields an infinite set of constraints on the allowed CFT data. Remarkably, these constraints can be turned into rigorous bounds without performing any path integral. The bootstrap thus provides exact, nonperturbative results about strongly coupled theories that are complementary to—and in some cases surpass—the results obtainable from lattice simulations, large N expansions, or perturbative RG.

The greatest success of the conformal bootstrap is its application to the three-dimensional Ising model. The 3d Ising CFT describes the critical point of the Ising universality class—the same universality class governing the liquid-gas transition, uniaxial ferromagnets, and binary fluid mixtures. While the critical exponents of this theory have been estimated by Monte Carlo simulations and perturbative RG for decades, the bootstrap has determined them to unprecedented precision, surpassing all other methods.

In this paper, we develop the formalism of conformal field theory in $d \geq 3$ dimensions and derive the conformal bootstrap. Section 2 reviews the necessary background: the conformal algebra and its representations, conformal correlators, radial quantization, unitarity bounds, the OPE, and conformal blocks. Section 3 presents the conformal bootstrap proper: we show how to convert crossing symmetry on OPE decompositions into bounds on operator dimensions, work through the analytic argument in detail, and present numerical results in $d = 2$ (chosen for computational tractability and the availability of closed-form conformal blocks), including a simple projected bound and an SDPB-based calculation of the Ising kink. Section 4 discusses future directions of study and open questions related to the conformal bootstrap, and Section 5 concludes.

2 Background: Conformal Field Theory

The Wilsonian renormalization group (RG) flow demonstrates that a broad variety of quantum field theories reach scale-invariant fixed points at very long or very short distances. These theories, known as *conformal field theories* (CFTs), are often interacting yet still exactly solvable due to their extreme symmetry constraints.

A *conformal transformation* is a diffeomorphism $x \to x'$ that rescales the metric by a local factor:

$$g'_{\mu\nu}(x') = \Omega(x)\, g_{\mu\nu}(x). \tag{1}$$

A CFT is a field theory which is invariant under these transformations. Generally we think of conformal transformations as active maps of the field operators on a fixed flat background. Such theories care about angles but not distances; the dynamics for a CFT are the same at all scales.

CFTs are so constrained that they can be analyzed via a variety of unique nonperturbative techniques, but understanding these requires the development of completely novel mathematical

formalism. (One might think a theory with greater symmetry constraints should be objectively more tractable to study—but since these constraints permit us to take on even harder problems, in practice the study of CFTs is just as complex as usual QFT, though in a different way!) In this section we will review some of the mathematics of CFTs with an eye toward our eventual study of the conformal bootstrap.

2.1 Assumptions

Throughout this discussion we will primarily treat scalar fields for simplicity. We will also take $d \geq 3$ to avoid complications with the infinite-dimensional Virasoro algebra of 2d CFTs, and we will consider Euclidean time. We work in natural units.

2.2 Conformal transformations

In any local quantum field theory, continuous symmetries are realized through conserved currents via the Ward identity. Given a conserved current j^μ with $\partial_\mu j^\mu = 0$, the corresponding charge

$$Q(\Sigma) = \int_\Sigma dS_\mu \, j^\mu(x) \tag{2}$$

integrated over a closed codimension-one surface Σ is independent of continuous deformations of Σ that do not cross operator insertions: by Stokes' theorem, $\int_\Sigma dS_\mu \, j^\mu - \int_{\Sigma'} dS_\mu \, j^\mu = \int_V d^d x \, \partial_\mu j^\mu = 0$ whenever the region V between Σ and Σ' contains no operators. In this sense, $Q(\Sigma)$ is a *topological surface operator*—it depends only on which operators lie inside Σ, not on the precise shape of Σ. This topological property of integrals over generators is effectively the meaning of a conservation law in the path integral formulation of QFT.

Spacetime symmetries are generated by the stress tensor $T^{\mu\nu}$. For a general infinitesimal coordinate transformation $x^\mu \to x^\mu + \epsilon^\mu(x)$, the associated current is $j^\mu = T^{\mu\nu}\epsilon_\nu$. Its divergence is

$$\partial_\mu(T^{\mu\nu}\epsilon_\nu) = (\partial_\mu T^{\mu\nu})\epsilon_\nu + \tfrac{1}{2}T^{\mu\nu}(\partial_\mu\epsilon_\nu + \partial_\nu\epsilon_\mu), \tag{3}$$

where we have used the symmetry of the stress tensor, $T^{\mu\nu} = T^{\nu\mu}$. Stress tensor conservation, $\partial_\mu T^{\mu\nu} = 0$, kills the first term on the RHS. Hence, the current $T^{\mu\nu}\epsilon_\nu$ is conserved for any Killing vector satisfying $\partial_\mu\epsilon_\nu + \partial_\nu\epsilon_\mu = 0$—that is, for translations ($\epsilon^\mu = a^\mu$, generating P_μ) and rotations ($\epsilon^\mu = \omega^\mu{}_\nu x^\nu$, generating $M_{\mu\nu}$). Thus the Poincaré group is generated by precisely the currents $T^{\mu\nu}\epsilon_\nu$ with ϵ_ν satisfying this Killing equation.

Now for conformal transformations, we can impose additional restrictions on the form of the stress tensor. In particular (and in brief), since $T^{\mu\nu}$ measures variations of the action under variations of the metric, we have that overall rescalings of the metric $g_{\mu\nu} \to \Omega^2 g_{\mu\nu}$ are mediated by the *trace* of the stress tensor, T^μ_μ. This is easy to see: under a Weyl rescaling $g_{\mu\nu} \to (1 + 2\sigma(x))g_{\mu\nu}$, the variation of the action is

$$\delta S = \tfrac{1}{2}\int d^d x \, \sqrt{g}\, T^{\mu\nu}\delta g_{\mu\nu} = \int d^d x \, \sqrt{g}\, \sigma(x)\, T^\mu{}_\mu,$$

so scale invariance of the action under arbitrary local $\sigma(x)$ requires $T^\mu{}_\mu = 0$.

A traceless stress tensor, then, corresponds to a theory with no inherent length scale. So we might well ask: for what $\epsilon^\mu(x)$ is the current $T^{\mu\nu}\epsilon_\nu$ conserved in the case that $T^{\mu\nu}$ is traceless?

Decomposing the symmetric tensor $\partial_\mu\epsilon_\nu + \partial_\nu\epsilon_\mu$ into its trace and traceless parts, the second term in the divergence becomes

$$\tfrac{1}{2}T^{\mu\nu}(\partial_\mu\epsilon_\nu + \partial_\nu\epsilon_\mu) = \tfrac{1}{2}T^{\mu\nu}\big(\partial_\mu\epsilon_\nu + \partial_\nu\epsilon_\mu - \tfrac{2}{d}(\partial\cdot\epsilon)\,\delta_{\mu\nu}\big) + \tfrac{1}{d}(\partial\cdot\epsilon)\,T^\mu{}_\mu. \tag{4}$$

If we require ϵ^μ to satisfy the *conformal Killing equation*

$$\partial_\mu\epsilon_\nu + \partial_\nu\epsilon_\mu = \frac{2}{d}(\partial\cdot\epsilon)\,\delta_{\mu\nu}, \tag{5}$$

the traceless part vanishes by definition, and the entire expression reduces to $\frac{1}{d}(\partial\cdot\epsilon)\,T^\mu{}_\mu$. This is a weaker condition than the Killing equation $\partial_\mu\epsilon_\nu + \partial_\nu\epsilon_\mu = 0$: it permits ϵ^μ to have a nonzero symmetric part, as long as that part is pure trace. The current is then conserved if and only if the stress tensor is traceless.

The conformal Killing equation admits more solutions than the Killing equation. In $d \geq 3$, the Killing solutions are translations ($\epsilon^\mu = a^\mu$) and rotations ($\epsilon^\mu = \omega^\mu{}_\nu x^\nu$), which have $\partial\cdot\epsilon = 0$. The conformal Killing equation additionally admits $\epsilon^\mu = \lambda x^\mu$ (with $\partial\cdot\epsilon = d\lambda$), generating *dilatations* D, and $\epsilon^\mu = b^\mu x^2 - 2x^\mu(b\cdot x)$ (with $\partial\cdot\epsilon = -2(d-1)\,b\cdot x$), generating *special conformal transformations* K_μ. The latter can be understood geometrically as the composition of an inversion, a translation, and another inversion. Together with translations and rotations, these exhaust the solutions of the conformal Killing equation in $d \geq 3$.

In summary, the independent solutions, generators, and finite forms of conformal transformations are:

Transformation	Generator	ϵ^μ	Finite form
Translations	P_μ	a^μ	$x^\mu \to x^\mu + a^\mu$
Rotations	$M_{\mu\nu}$	$\omega^\mu{}_\nu x^\nu$	$x^\mu \to R^\mu{}_\nu x^\nu$
Dilatations	D	λx^μ	$x^\mu \to e^\lambda x^\mu$
Special conformal	K_μ	$b^\mu x^2 - 2x^\mu(b\cdot x)$	$x^\mu \to \dfrac{x^\mu - b^\mu x^2}{1 - 2b\cdot x + b^2x^2}$

2.3 The conformal algebra, primaries, and descendants

The operators described in the last section generate the conformal group, which in $d \geq 3$ turns out to be isomorphic to $SO(d+1,1)$. (In $d = 2$, the conformal group is generated by the infinite-dimensional Virasoro algebra; we will ignore this case in this section.) In Euclidean signature, the non-vanishing

commutators of the conformal algebra are:

$$[D,\, P_\mu] = P_\mu, \tag{6}$$
$$[D,\, K_\mu] = -K_\mu, \tag{7}$$
$$[K_\mu,\, P_\nu] = 2\big(\delta_{\mu\nu}\, D - M_{\mu\nu}\big), \tag{8}$$
$$[M_{\mu\nu},\, P_\rho] = \delta_{\mu\rho}\, P_\nu - \delta_{\nu\rho}\, P_\mu, \tag{9}$$
$$[M_{\mu\nu},\, K_\rho] = \delta_{\mu\rho}\, K_\nu - \delta_{\nu\rho}\, K_\mu, \tag{10}$$
$$[M_{\mu\nu},\, M_{\rho\sigma}] = \delta_{\mu\rho}\, M_{\nu\sigma} - \delta_{\mu\sigma}\, M_{\nu\rho} - \delta_{\nu\rho}\, M_{\mu\sigma} + \delta_{\nu\sigma}\, M_{\mu\rho}. \tag{11}$$

To obtain representations of this algebra we need to characterize the action of the Poincaré, dilatation, and special conformal generators on operators. Poincaré representations are characterized by translation and rotation quantum numbers in the usual way: translations act by

$$\mathcal{O}(x) = e^{-x^\mu P_\mu}\, \mathcal{O}(0)\, e^{x^\mu P_\mu} \tag{12}$$

and rotations act on operators at the origin via a representation of $SO(d)$, characterized by the spin ℓ. We will focus primarily on scalar operators in what follows, noting that the generalization to spinning operators is straightforward but notationally heavier.

The dilatation generator D extends the Poincaré algebra to include scale transformations. An operator $\mathcal{O}$ at the origin is said to have *scaling dimension* Δ if under a finite dilatation $x \to \lambda x$ it transforms as

$$\mathcal{O}(x) \;\to\; \lambda^\Delta\, \mathcal{O}(\lambda x). \tag{13}$$

That such eigenstates always exist follows from the structure of radial quantization, discussed below, through which we will view the dilatation operator as a time evolution operator with a set of orthonormal eigenstates.

Note from the commutators $[D,\, P_\mu] = P_\mu$ and $[D,\, K_\mu] = -K_\mu$ that P_μ and K_μ are raising and lowering operators for the scaling dimension Δ: if $\mathcal{O}$ has dimension Δ, then

$$\Delta(P_\mu\, \mathcal{O}) = \Delta + 1, \qquad \Delta(K_\mu\, \mathcal{O}) = \Delta - 1. \tag{14}$$

Analogously to the usual argument for raising and lowering operators from canonical quantization, we have that since dimensions must be bounded below for a sensible theory, there must be some operators with lowest possible scaling dimension such that

$$[K_\mu,\, \mathcal{O}(0)] = 0. \tag{15}$$

These *primary* operators provide lowest-weight elements from which we can construct all other higher-dimension operators in a representation, known as *descendants*.

In brief, any primary operator leads to a representation of the conformal algebra via applications of P_μ. This is analogous to the development of irreducible representations of $SU(2)$ from highest-weight states.

We will find in our discussion of reflection positivity that any local operator in a unitary CFT (or its Wick transform) is a linear combination of primaries and descendants.

2.4 Conformal Correlators

2.4.1 The 2-point function

Correlation functions are extremely constrained by conformal invariance. This is readily apparent in the two-point function: via translational, rotational, and dilatation symmetry we can place one of the two operators at the origin, rotate the other onto (say) the x-axis, and scale the function down until the two operators are one unit apart, placed at $x = 0$ and $x = 1$ respectively. The value of the correlator can only depend on the scaling factors obtained from these transformations (up to a constant). Working through this explicitly, translations, rotations, and dilatations constrain the two-point function to the form

$$\langle \mathcal{O}_i(x)\, \mathcal{O}_j(0)\rangle = \frac{c_{ij}}{|x|^{\Delta_i+\Delta_j}}, \tag{16}$$

where c_{ij} is an undetermined constant. We can actually constrain this even further with special conformal transformations (SCTs). Under an SCT with parameter b^μ, a primary of dimension Δ at position x picks up a prefactor $|1 - 2b\cdot x + b^2x^2|^\Delta$. Plugging this into Equation 16 and demanding invariance of the two point function under SCTs gives

$$|1 - 2b\cdot x + b^2x^2|^{(\Delta_j-\Delta_i)/2} = 1 \tag{17}$$

for all b^μ and x, which forces $\Delta_i = \Delta_j$. Additionally, we will find in our later discussion of reflection positivity that we can choose an orthonormal basis of primary operators with trivial coefficients in their two point functions. For elements of such a basis, then, the two-point function takes the simple form

$$\langle \mathcal{O}_i(x)\, \mathcal{O}_j(0)\rangle = \frac{\delta_{ij}}{|x|^{2\Delta_i}}. \tag{18}$$

2.4.2 The 3-point function

Now consider a 3-point function; for similar reasons as above, we can constrain such a correlator completely via conformal symmetry. We have already seen that we have enough freedom to place one operator at the origin and the other one unit away on the x-axis. Special conformal transformations can then be used to place the third operator at infinity. This procedure gives the 3-point correlator

$$\langle \mathcal{O}_1(x_1)\, \mathcal{O}_2(x_2)\, \mathcal{O}_3(x_3)\rangle = \frac{f_{123}}{|x_{12}|^{\Delta_1+\Delta_2-\Delta_3}\ |x_{13}|^{\Delta_1+\Delta_3-\Delta_2}\ |x_{23}|^{\Delta_2+\Delta_3-\Delta_1}}, \tag{19}$$

where $x_{ij} \equiv x_i - x_j$ and the *OPE coefficient* f_{123} is the sole piece of dynamical data not fixed by symmetry. (We will relate this number to operator product expansions later.) The exponents are determined by requiring that the correlator scale correctly under $x \to \lambda x$ for each operator individually. For scalar primaries, the three-point function is therefore completely specified up to the single number f_{123}.

2.4.3 The 4-point function

Constraining the 3-point functions has utilized nearly all of the different conformal generators at our disposal, with the exception of the generator of rotations around the x-axis. If we add another operator, all we can do is rotate it so it lies in (for example) the x, y plane; we cannot

fix its position entirely. Hence the 4-point function depends on two independent *cross ratios* in the correlator—roughly the x and y coordinates of the final fourth operator after applying the transformations described above.

Working through this explicitly for four identical scalar primaries ϕ with dimension Δ_ϕ, conformal invariance constrains the four-point function to the form

$$\langle \phi(x_1)\,\phi(x_2)\,\phi(x_3)\,\phi(x_4)\rangle = \frac{g(u,v)}{|x_{12}|^{2\Delta_\phi}\,|x_{34}|^{2\Delta_\phi}}, \tag{20}$$

where $g(u,v)$ is an arbitrary function of the two conformally invariant cross ratios

$$u = \frac{x_{12}^2\,x_{34}^2}{x_{13}^2\,x_{24}^2}, \qquad v = \frac{x_{14}^2\,x_{23}^2}{x_{13}^2\,x_{24}^2}. \tag{21}$$

Since the four ϕ's are identical, the correlator must be invariant under permutations of the operator insertion locations. Invariance under $1 \leftrightarrow 3$ (which sends $u \leftrightarrow v$) imposes the *crossing symmetry* condition

$$g(u,v) = \left(\frac{u}{v}\right)^{\Delta_\phi} g(v,u). \tag{22}$$

Similarly, $1 \leftrightarrow 2$ (which sends $u \to u/v$, $v \to 1/v$) gives an equivalent constraint. These conditions, which arise purely from the consistency of the OPE, will be the central ingredient of the conformal bootstrap.

2.5 Radial Quantization and the State-Operator Correspondence

In the path integral formulation of quantum field theory, choosing a quantization scheme is tantamount to choosing how to foliate spacetime into equal-time slices: one performs the path integral over fields on one side of a slice to produce a state on that slice. The standard choice in flat Euclidean space is to foliate by parallel hyperplanes $x^0 = \text{const}$, but nothing prevents us from choosing a different foliation.

In *radial quantization*, we foliate $\mathbb{R}^d$ by concentric spheres S^{d-1} centered at the origin. A "time slice" is now a sphere of radius r; the path integral over the interior of the sphere produces a state $|\psi\rangle$ on that sphere. Time evolution—propagation from one sphere to a larger one—is generated by the dilatation operator D, which plays the role of the Hamiltonian. Explicitly, the map $r \to e^\tau$ identifies radial quantization on $\mathbb{R}^d$ with ordinary quantization on the cylinder $\mathbb{R}_\tau \times S^{d-1}$, where $\tau = \log r$ is the cylinder time.

The key consequence of this construction is the *state-operator correspondence*: there is a one-to-one map between local operators at the origin and states on S^{d-1}. Getting from operators to states is trivial: just insert some operator $\mathcal{O}(0)$ at the origin and (radially) time evolve out to a sphere of a given radius (say 1) to obtain a state $\mathcal{O}(0)\,|0\rangle$. In this sense we can think of the ground state as dual to the identity operator. From here we can radially evolve this state out to any radius.

To get from states to operators is also fairly simple. We can start with an eigenstate $|\mathcal{O}_i\rangle$ of the dilatation operator

$$D\,|\mathcal{O}_i\rangle = \Delta_i\,|\mathcal{O}_i\rangle$$

defined on a sphere of radius 1, and time evolve this state backwards toward the origin. In the

limit, the state becomes an infinitesimally small ball at the origin that can be inserted into path integrals and evolved outwards—exactly what a local operator is. Hence we have the state operator correspondence:

$$\mathcal{O}(0) \leftrightarrow \mathcal{O}(0)\,|0\rangle \equiv |\mathcal{O}\rangle\,.$$

This relation is the backbone of the bootstrap, as we will see.

2.6 Reflection Positivity and Unitarity Bounds

In Lorentzian quantum mechanics, unitarity requires that the Hilbert space have a positive-definite inner product: $\langle\psi|\psi\rangle \geq 0$ for all states, with equality only for the zero state. The Euclidean counterpart of this condition is *reflection positivity*. Given a reflection θ about some codimension-one surface (say the hyperplane $x^0 = 0$), reflection positivity requires that for any operator $\mathcal{O}$ supported in the half-space $x^0 > 0$,

$$\langle(\theta \cdot \mathcal{O})^\dagger\, \mathcal{O}\rangle \geq 0. \tag{23}$$

In radial quantization, where the relevant reflection is the inversion $x^\mu \to x^\mu/|x|^2$, this is precisely the statement that the norm of the state $|\mathcal{O}\rangle$ created by inserting $\mathcal{O}$ at the origin is non-negative. Reflection positivity in the Euclidean theory is therefore equivalent to unitarity (positive-definite Hilbert space) in the Lorentzian theory. We can view this intuitively as a requirement on the positivity of an inner product of two states, each obtained by a path integral up to the given boundary from opposite directions with reflected operator insertions.

2.6.1 Unitarity bounds

This positivity condition places powerful constraints on the allowed operator dimensions, known as *unitarity bounds*. The idea is to consider descendant states and demand that they have non-negative norm. In radial quantization $K_\mu = P_\mu^\dagger$, so the norm of the level-1 descendant $P_\mu|\mathcal{O}\rangle$ is

$$\|P_\mu|\mathcal{O}\rangle\|^2 = \langle\mathcal{O}|K_\mu P_\mu|\mathcal{O}\rangle. \tag{24}$$

Since $K_\mu|\mathcal{O}\rangle = 0$ for a primary, we can replace $K_\mu P_\mu$ by the commutator $[K_\mu, P_\mu]$. Using $[K_\mu,\, P_\nu] = 2(\delta_{\mu\nu}\, D - M_{\mu\nu})$ and tracing over μ:

$$[K^\mu,\, P_\mu] = 2\big(d\, D - M^\mu{}_\mu\big) = 2d\, D, \tag{25}$$

since $M^\mu{}_\mu = 0$ by antisymmetry. On a scalar primary ($M_{\mu\nu}|\mathcal{O}\rangle = 0$, $D|\mathcal{O}\rangle = \Delta|\mathcal{O}\rangle$) this gives

$$\sum_\mu \|P_\mu|\mathcal{O}\rangle\|^2 = 2d\,\Delta\,\langle\mathcal{O}|\mathcal{O}\rangle \geq 0, \tag{26}$$

which requires $\Delta \geq 0$, with equality only for the identity operator.

This is the simplest unitarity bound, coming from positive-definiteness of the inner product of level-1 descendants. It turns out the strongest bounds we can obtain require consideration of just level-2 descendants; for brevity we will simply state the result:

$$\Delta \geq \begin{cases} \frac{d-2}{2} & \ell = 0, \\ d + \ell - 2 & \ell \geq 1, \end{cases} \tag{27}$$

with the identity operator ($\Delta = \ell = 0$) as the unique exception. These bounds are essential to the bootstrap: they restrict the domain of allowed CFT data over which the crossing equations must be satisfied.

2.6.2 Primaries and descendants span the Hilbert space

Finally, reflection positivity guarantees D is self-adjoint, and thus guarantees that the Hilbert space in radial quantization is spanned by states of definite dimension and spin—i.e., that D is diagonalizable with real eigenvalues. Combined with the state-operator correspondence, this means that every local operator in a unitary CFT can be written as a linear combination of primaries and their descendants. There are no other operators; the primary-descendant decomposition is complete. We used this fact when constraining the two-point function, and it will be important in what follows.

2.7 The Operator Product Expansion

The preceding ingredients are sufficient to naturally derive the operator product expansion. Apply two operators $\mathcal{O}_i(x)\mathcal{O}_j(0)$ inside a ball centered at the origin. Perform the path integral over the interior of the ball; this gives a state on the boundary. But any state is dual to a linear combination of primaries and descendants at the origin; hence, we can decompose this product of operators as

$$\mathcal{O}_i(x)\,\mathcal{O}_j(0) = \sum_k \frac{f'_{ijk}}{|x|^{\Delta_i+\Delta_j-\Delta_k}}\, C_{ijk}(x,\,\partial)\,\mathcal{O}_k(0), \tag{28}$$

where the sum runs over all primary operators $\mathcal{O}_k$ in the theory, f'_{ijk} are parameters that we will soon identify with the OPE coefficients f_{ijk} from the 3 point function (19), and $C_{ijk}(x,\,\partial)$ are differential operators that encode the contributions of all descendants of $\mathcal{O}_k$. Note this is normalized so that the leading term in C_{ijk} is 1. For example, for identical external scalars ϕ exchanging a scalar primary $\mathcal{O}$ of dimension Δ,

$$C_{\phi\phi\mathcal{O}}(x,\partial) = 1 + \frac{1}{2}x^\mu\partial_\mu + O(\partial^2), \tag{29}$$

where the coefficient $\frac{1}{2}$ of the first descendant contribution is universal for identical scalars and the higher-order terms are fully determined by Δ, Δ_ϕ, and d.

To see that the coefficients f'_{ijk} are exactly the OPE coefficients f_{ijk} that appear in the three-point function (19), act on both sides of (28) with a third operator $\mathcal{O}_l(y)$ at a well-separated point and take the expectation value. On the left-hand side, we get exactly a three-point function whose value we have already constrained to be

$$\langle \mathcal{O}_i(x)\,\mathcal{O}_j(0)\,\mathcal{O}_l(y)\rangle = \frac{f_{ijl}}{|x|^{\Delta_i+\Delta_j-\Delta_l}\,|y|^{\Delta_j+\Delta_l-\Delta_i}\,|x-y|^{\Delta_i+\Delta_l-\Delta_j}}.$$

On the right-hand side, orthonormality of the two-point function (18) kills all terms in the sum other than $k = l$, and orthogonality of a primary with its descendants kills everything in C_{ijl} past

the leading order term (which is 1). Matching the two sides and using the expression for the 2-point function (18) identifies $f'_{ijk} = f_{ijk}$.

This argument also shows that the differential operators C_{ijk} are completely fixed by conformal symmetry: they are determined by requiring that the OPE reproduces the correct conformal transformation properties of all descendant contributions, order by order in derivatives. Once the leading coefficient is normalized to unity, the rest of C_{ijk} is a purely kinematic object encoding the descendant structure of the conformal multiplet. Hence the only variable parameters determining the structure of the theory are the coefficients f_{ijk}.

In an orthonormal basis of operators, the OPE coefficients f_{ijk} are real. This follows from reflection positivity when we take the inner product of both sides of (28) with some operator $\mathcal{O}_k$. We conclude $f^2_{\phi\phi\mathcal{O}} \geq 0$, which will be crucial to what follows.

2.8 Conformal Blocks

The power of the OPE becomes apparent when applied to four-point functions. Consider the correlator of four identical scalars $\langle \phi(x_1)\, \phi(x_2)\, \phi(x_3)\, \phi(x_4) \rangle$. We apply the OPE to the pair $\phi(x_1)\, \phi(x_2)$, expanding in primaries $\mathcal{O}_k$ and their descendants at x_2, and similarly to $\phi(x_3)\, \phi(x_4)$ at x_4. The four-point function reduces to a sum over two-point functions of $\mathcal{O}_k$ and its descendants, each of which is fixed by conformal symmetry.

Explicitly, inserting the OPE (28) on both pairs gives

$$\langle \phi(x_1)\phi(x_2)\phi(x_3)\phi(x_4) \rangle = \sum_{k,k'} \frac{f_{\phi\phi k}\, f_{\phi\phi k'}}{|x_{12}|^{2\Delta_\phi - \Delta_k}\, |x_{34}|^{2\Delta_\phi - \Delta_{k'}}}\, C_{\phi\phi k}(x_{12}, \partial_2)\, C_{\phi\phi k'}(x_{34}, \partial_4)\, \langle \mathcal{O}_k(x_2)\, \mathcal{O}_{k'}(x_4) \rangle. \tag{30}$$

The two-point function vanishes unless $k = k'$ by Equation 18, so the double sum collapses to a single sum over primaries $\mathcal{O}$ exchanged in the $(12) \to (34)$ channel. For each exchanged primary, the action of the differential operators $C_{\phi\phi\mathcal{O}}$ on the two-point function, summed over all descendants of $\mathcal{O}$, produces a single function of the cross ratios—this function is the conformal block $g_{\Delta,\ell}(u, v)$. Hence we have:

$$\langle \phi(x_1)\phi(x_2)\phi(x_3)\phi(x_4) \rangle = \frac{1}{|x_{12}|^{2\Delta_\phi}\, |x_{34}|^{2\Delta_\phi}} \sum_{\mathcal{O}} f^2_{\phi\phi\mathcal{O}}\, g_{\Delta,\ell}(u, v). \tag{31}$$

Notice that (for this case where operators are identical) the same OPE coefficient $f_{\phi\phi\mathcal{O}}$ appears in both the (12) and (34) channels, which is why Equation 31 contains $f^2_{\phi\phi\mathcal{O}}$ rather than two independent coefficients.

Using our previous expression for the 4-point function, Equation 20, and stripping off the kinematic prefactor in Equation 31, this gives the *conformal block decomposition*:

$$g(u, v) = 1 + \sum_{\mathcal{O} \neq \mathbf{1}} f^2_{\phi\phi\mathcal{O}}\, g_{\Delta,\ell}(u, v), \tag{32}$$

where the sum runs over all primaries $\mathcal{O}$ of dimension Δ and spin ℓ appearing in the $\phi \times \phi$ OPE (only even spins contribute, by Bose symmetry of identical scalars). We have separated the identity operator contribution, which gives the 1. The functions $g_{\Delta,\ell}(u, v)$ are the *conformal blocks*: each one encodes the total contribution to the four-point function from a single primary $\mathcal{O}$ and all of its

descendants. They are purely kinematic objects, determined entirely by conformal symmetry given the exchanged quantum numbers (Δ, ℓ) and the external dimension Δ_ϕ.

2.8.1 Computing conformal blocks

The last ingredient we need in order to discuss the bootstrap is an analytic and/or computationally efficient way to compute the value of a conformal block. We keep this brief; for a more thorough discussion see [4], Chapter 9.3.

Conformal blocks can be thought of as the contributions to the 4-point function from different irreducible representations of the conformal group. The Casimir operator C for the conformal group acts with the same eigenvalue on every state in a given irrep. As a consequence of this, every conformal block $g_{\Delta,\ell}(u,v)$ is an eigenfunction of the Casimir with eigenvalue $\lambda_{\Delta,\ell}$. Since the Casimir acts as a differential operator, we find that conformal blocks satisfy a differential equation:

$$\mathcal{D} g_{\Delta,\ell}(u,v) = \lambda_{\Delta,\ell} g_{\Delta,\ell}(u,v)$$

where $\mathcal{D}$ is the differential operator giving the action of the Casimir in the representation corresponding to $g_{\Delta,\ell}$ and at cross ratios u, v. This *Casimir equation*, together with the boundary condition that $g_{\Delta,\ell}(u,v) \sim u^{(\Delta-\ell)/2}$ as $u \to 0$ (corresponding to the leading OPE singularity), uniquely determines the conformal block.

In even spacetime dimensions, the Casimir equation can be solved analytically. In general (odd) d, no such closed form is known, and one must resort to series expansions—typically a power series in $\sqrt{u}$ (the "radial expansion") whose coefficients are determined recursively from the Casimir equation. These expansions converge rapidly and are the basis for the numerical bootstrap in odd dimensions.

3 The Conformal Bootstrap

3.1 OPE Associativity and Crossing Symmetry

We have established that a CFT is completely characterized by its CFT data: the spectrum of primary operators Δ_i, ℓ_i and the OPE coefficients f_{ijk}. Using the OPE, any n-point correlation function can be reduced to $(n-1)$-point functions, and by iterating, ultimately to the known two-point functions. All correlation functions are therefore determined by the CFT data.

But given an arbitrary set of numbers Δ_i, f_{ijk}, do they always necessarily define a consistent CFT? The answer is no: the OPE must be *associative*. When we reduce a four-point function using the OPE, we can choose to pair the operators in different ways. For instance, given $\langle \mathcal{O}_1 \mathcal{O}_2 \mathcal{O}_3 \mathcal{O}_4 \rangle$, we can apply the OPE to the pair $(\mathcal{O}_1, \mathcal{O}_2)$ and to $(\mathcal{O}_3, \mathcal{O}_4)$—the (12) → (34) or *s-channel*—or instead to $(\mathcal{O}_1, \mathcal{O}_4)$ and $(\mathcal{O}_2, \mathcal{O}_3)$—the (14) → (23) or *t-channel*. Both must give the same result.

In either channel, each exchanged primary $\mathcal{O}_k$ contributes a conformal block determined by its dimension Δ_k and spin ℓ_k, and the block depends on the cross ratios u, v formed from the four insertion points. Thus, denoting the s- and t-channel blocks by $g^{(s)}_{\Delta_k,\ell_k}(u,v)$ and $g^{(t)}_{\Delta_k,\ell_k}(u,v)$ respectively, associativity of the OPE requires

$$\sum_k f_{12k} f_{34k} g^{(s)}_{\Delta_k,\ell_k}(u,v) = \sum_k f_{14k} f_{23k} g^{(t)}_{\Delta_k,\ell_k}(u,v). \tag{33}$$

The two sums run over the *same* set of primaries—those appearing in both $\mathcal{O}_1 \times \mathcal{O}_2$ and $\mathcal{O}_3 \times \mathcal{O}_4$ (and equivalently in $\mathcal{O}_1 \times \mathcal{O}_4$ and $\mathcal{O}_2 \times \mathcal{O}_3$)—but with manifestly different OPE coefficients and different kinematic dependence on the cross ratios u, v. In particular, the two-channel blocks $g^{(s)}$ and $g^{(t)}$ are related by the permutation of insertion points that swaps $1 \leftrightarrow 3$; for identical external scalars this is an especially simple relation, given by the $u \leftrightarrow v$ exchange we encountered in Equation 22. Equation 33 is the *crossing symmetry equation*: the requirement that the conformal block decomposition is independent of the channel in which we apply the OPE.

Note also that crossing symmetry is a constraint on exactly the same data that appears in the three-point functions. The OPE coefficients f_{ijk} are the dynamical content of both the three-point function Equation 19 and the conformal block decomposition of the four-point function. Crossing symmetry therefore links different three-point functions to one another: the coefficient f_{12k} that appears in $\langle \mathcal{O}_1 \mathcal{O}_2 \mathcal{O}_k \rangle$ is constrained by its appearance alongside f_{34k} from $\langle \mathcal{O}_3 \mathcal{O}_4 \mathcal{O}_k \rangle$ in the crossing equation.

3.2 The Crossing Equation for Identical Scalars

We now specialize to the simplest and most-studied case: a four-point function of identical real scalars ϕ with dimension Δ_ϕ. Here, the conformal blocks $g^{(s)}_{\Delta_k,\ell_k}(u,v)$ and $g^{(t)}_{\Delta_k,\ell_k}(u,v)$ are related simply by the crossing condition, Equation 22, which tells us exactly how the four-point function changes under permutation of operator insertion coordinates. Applying this to the conformal block decomposition Equation 32, we have

$$\sum_{\mathcal{O}} f^2_{\phi\phi\mathcal{O}}\, g_{\Delta,\ell}(u,v) = \left(\frac{u}{v}\right)^{\Delta_\phi} \sum_{\mathcal{O}} f^2_{\phi\phi\mathcal{O}}\, g_{\Delta,\ell}(v,u). \tag{34}$$

Here the sums run over all primaries (including the identity) in the $\phi \times \phi$ OPE with even spin. Rearranging, we obtain the crossing equation in the form that is most useful for the bootstrap:

$$\sum_{\mathcal{O}} f^2_{\phi\phi\mathcal{O}}\, F^{\Delta_\phi}_{\Delta,\ell}(u,v) = 0, \tag{35}$$

where

$$F^{\Delta_\phi}_{\Delta,\ell}(u,v) \equiv v^{\Delta_\phi}\, g_{\Delta,\ell}(u,v) - u^{\Delta_\phi}\, g_{\Delta,\ell}(v,u). \tag{36}$$

Equation 35 is the master equation of the conformal bootstrap. It says that a specific set of functions $F^{\Delta_\phi}_{\Delta,\ell}$—one for each primary in the OPE—must sum to zero with positive coefficients $f^2_{\phi\phi\mathcal{O}} \geq 0$. This is a highly nontrivial constraint: generically, a collection of functions cannot sum to zero with all-positive coefficients. Note we know what the functions $F^{\Delta_\phi}_{\Delta,\ell}(u,v)$ are (at least approximately) thanks to the Casimir equation.

It is convenient to introduce the variables $z, \bar{z}$ defined by

$$u = z\bar{z}, \qquad v = (1-z)(1-\bar{z}). \tag{37}$$

The crossing symmetry $u \leftrightarrow v$ corresponds to $(z,\bar{z}) \to (1-z, 1-\bar{z})$. The crossing-symmetric point is $z = \bar{z} = \frac{1}{2}$, corresponding to $u = v = \frac{1}{4}$.

3.3 Bounding the space of CFTs

The crossing equation (35) can be used to derive rigorous bounds on the CFT data. The key insight, due to Rattazzi, Rychkov, Tonni, and Vichi [3], is the following. We can think of each function $F^{\Delta_\phi}_{\Delta,\ell}(u,v)$ as a vector $\vec{F}_{\Delta,\ell}$ in an infinite-dimensional function space. The crossing equation then says that these vectors sum to zero with positive coefficients:

$$\sum_{\Delta,\ell} p_{\Delta,\ell}\,\vec{F}_{\Delta,\ell} = 0, \qquad p_{\Delta,\ell} = f^2_{\phi\phi\mathcal{O}} \geq 0. \tag{38}$$

This is highly restrictive. If we consider any hyperplane through the origin in the function space of $\vec{F}_{\Delta,\ell}$, in order for Equation 35 to be satisfied there must be functions in the OPE decomposition of $\phi \times \phi$ that lie on either side of this hyperplane; if they all lie strictly on one side then it is impossible to sum them to zero with positive coefficients.

Hence given some hypothesis about the space of functions in the OPE, we search for a *linear functional* α, acting on functions of u, v, such that

$$\alpha(F^{\Delta_\phi}_{\Delta,\ell}) \geq 0 \tag{39}$$

for all functions satisfying the unitarity bounds as well as our hypothesis, and with strict positivity on at least one operator (typically the identity, whose coefficient $f^2_{\phi\phi\mathbf{1}} = 1$ is fixed). If such a functional exists, we can apply α to both sides of Equation 35: the left-hand side is a sum of non-negative terms with at least one strictly positive, giving a positive number, while the right-hand side is zero—a contradiction. The hypothesis is therefore ruled out.

This leads, for example, to the following algorithm for bounding operator dimensions:

1. Hypothesize that no scalar in the $\phi \times \phi$ OPE has dimension below Δ_*.
2. Search for a linear functional α satisfying Equation 39 for all $\Delta \geq \Delta_*$ (scalars), $\Delta \geq d + \ell - 2$ (spinning operators), and $\alpha(F^{\Delta_\phi}_{0,0}) > 0$ (identity).
3. If α exists, the hypothesis is inconsistent: some scalar must have $\Delta < \Delta_*$.
4. Increase Δ_* and repeat. The largest Δ_* for which the hypothesis can be excluded gives an upper bound on the dimension of the lowest-lying scalar in the OPE.

An especially useful fact that makes this computation actually tractable is that we can project the entire algorithm onto a subspace of the full function space; for example we could work in the 2-dimensional subspace corresponding to evaluations of the function at two different values of u, v. The more components/function evaluations we add, of course, the stronger the restrictions will be.

3.4 An Example Bound

To illustrate the method concretely, we work through an example calculation analogous to that of chapter 10.5 in [4], but with a different numerical setup. We consider a 2d CFT with a real scalar primary ϕ of dimension

$$\Delta_\phi = 0.20. \tag{40}$$

As mentioned previously, in $d = 2$, the conformal blocks for identical external scalars are known in closed form:

$$g_{\Delta,\ell}(z,\bar{z}) = k_{\Delta+\ell}(z)\, k_{\Delta-\ell}(\bar{z}) + k_{\Delta+\ell}(\bar{z})\, k_{\Delta-\ell}(z), \tag{41}$$

where $k_\beta(z) = z^{\beta/2}\, {}_2F_1(\frac{\beta}{2}, \frac{\beta}{2}; \beta;\, z)$.

To reduce the problem to a tractable finite-dimensional one, we project the crossing equation onto a two-dimensional subspace. Define

$$\vec{v}(F) = (H(0.5, 0.65) - H(0.5, 0.30),\ H(0.45, 0.60) - H(0.30, 0.25)) \in \mathbb{R}^2, \tag{42}$$

where $H(z,\bar{z}) = F(u,v)/(u^{\Delta_\phi} - v^{\Delta_\phi})$, with $u = z\bar{z}$ and $v = (1-z)(1-\bar{z})$. Since the crossing equation $\sum p_{\Delta,\ell}\, F_{\Delta,\ell} = 0$ holds as an equation of functions, it also holds after applying the projection $\vec{v}$:

$$\sum_{\Delta,\ell} p_{\Delta,\ell}\, \vec{v}(F^{\Delta_\phi}_{\Delta,\ell}) = 0. \tag{43}$$

We calculate the values of the projected components $\vec{v}(F^{\Delta_\phi}_{\Delta,\ell})$ for even spins up to $\ell = 10$ and 5000 different values of Δ from $\Delta = 0$ to $\Delta = 8$ for scalars (with finer spacing near the unitarity bound) and from $\Delta = \ell$ to $\Delta = \ell + 30$ for each even spin ℓ. We start our sampling for values of Δ saturating the unitarity bounds.

Figure 1 shows the resulting projected curves. As in the standard geometric argument, the dashed line through the origin is a candidate separating hyperplane (functional $\alpha = \vec{u} \cdot \vec{v}$), chosen using the stress-tensor direction at $\ell = 2$. The highlighted segment of the scalar curve is the interval that lies on the opposite side of this line from the spinning-sector curves. Evidently, in order for a positive-coefficient linear combination of these function values to equal zero, there must be some scalar in the OPE with scaling dimension Δ in the highlighted region opposite the dashed line from the nonzero spin components.

Hence we find that, in a unitary 2d CFT with a real scalar operator ϕ of scaling dimension $\Delta_\phi = 0.2$, there must be at least one scalar in the $\phi \times \phi$ OPE with dimension $\Delta \in [0.2564,\ 1.3463]$. Other projections with different parameter choices are collected in Appendix A.

3.5 Numerical Bootstrap Techniques

The example above used a two-dimensional projection of the crossing equation, and the resulting bound was already nontrivial. By working in higher-dimensional subspaces, one can obtain much stronger bounds. In practice, one parametrizes the linear functional as

$$\alpha(F) = \sum_{m+n \le \Lambda} a_{mn}\, \partial_z^m\, \partial_{\bar{z}}^n\, F(z,\bar{z})\big|_{z=\bar{z}=\frac{1}{2}}, \tag{44}$$

where Λ is a cutoff that controls the dimension of the search space. Evaluating derivatives at $z = \bar{z} = \frac{1}{2}$ in this way turns out to be more numerically efficient than working with linear combinations of components, since the crossing equation and conformal blocks converge exponentially fast at that point.

There are two central computational challenges in implementing this program at scale. First, the space of allowed functionals α is finite-dimensional (parametrized by the $\frac{1}{2}\lfloor\frac{\Lambda+1}{2}\rfloor(\lfloor\frac{\Lambda+1}{2}\rfloor + 1)$ coefficients a_{mn}), but the positivity constraints $\alpha(F^{\Delta_\phi}_{\Delta,\ell}) \ge 0$ must hold for a continuous range of Δ.

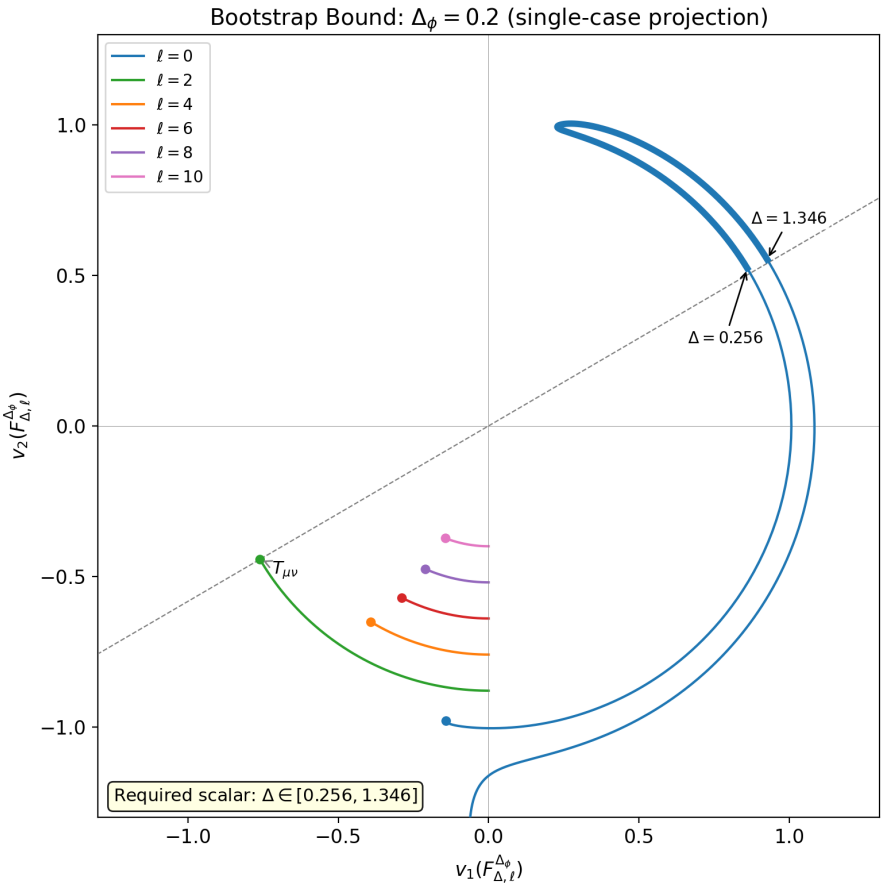

Figure 1: Two-dimensional projection of crossing vectors for $\Delta_\phi = 0.20$ using $\vec{v}(F) = \big(H(0.5, 0.65) - H(0.5, 0.30),\ H(0.45, 0.60) - H(0.30, 0.25)\big)$. The thick blue segment marks the scalar dimensions excluded from the same half-plane as the spinning curves. Solid dots represent values of Δ that saturate the unitarity bounds for that value of ℓ.

Several approaches have been used to handle this. The original work of Rattazzi et al. [3] discretized Δ on a grid and solved the resulting finite system of linear inequalities using linear programming. A more refined approach, developed by El-Showk and Paulos [6], uses a customized version of the simplex algorithm that handles the continuously varying constraints directly. The best method now is to reformulate the positivity constraints as conditions on polynomials and solve the resulting problem using *semidefinite programming* (SDP). Since the conformal blocks are rational functions of Δ (when expanded in the radial coordinate), the positivity constraints can be written as polynomial matrix inequalities, which are precisely the constraints that SDP is designed to optimize over. The software package `SDPB`, developed by Simmons-Duffin [7], implements a specialized SDP solver tailored to the structure of bootstrap problems and has enabled the highest-precision results to date.

Second, efficient evaluation of conformal blocks in general spacetime dimension is nontrivial. In even d, the blocks have closed-form expressions in terms of hypergeometric functions (as we noted above), but in odd d—including the physically important case $d = 3$—one must use series expansions or recursion relations. The Zamolodchikov-type recursion relation [8] expresses conformal blocks in terms of lower-dimensional data and converges rapidly; this provides the numerical infrastructure

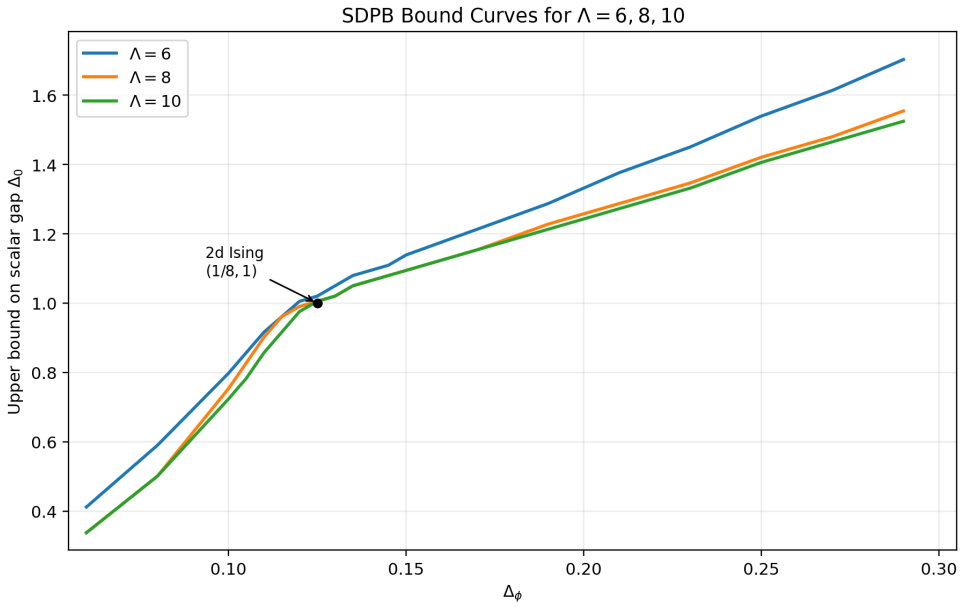

Figure 2: SDPB bound curves generated in this project for $\Lambda = 6, 8, 10$ using the Section 10.7 procedure of [4]. The scan uses 19 values of Δ_ϕ, with increased density near $\Delta_\phi = \frac{1}{8}$, and the kink appears near the 2d Ising point $(\Delta_\phi, \Delta_0) = (\frac{1}{8}, 1)$.

on which modern bootstrap computations are built.

3.6 The Ising Model

The procedure in Section 3.4 can be repeated for any value of Δ_ϕ. A particularly interesting case is $\Delta_\phi = \frac{1}{8}$, which is the known dimension of the spin operator in the 2d Ising model. Using the same basic two-point projection as in our example, Simmons-Duffin [4] obtains the bound $\Delta_\epsilon \in [0.161, 1.04]$. The exact 2d Ising value is $\Delta_\epsilon = 1$, so even this simple analysis places the upper bound within 4% of the correct answer—a nontrivial constraint from nothing more than unitarity and crossing symmetry evaluated at two points.

Applying the more powerful numerical techniques of the preceding section dramatically sharpens this bound. We carried out a basic SDPB implementation, following the Section 10.7 workflow of [4]: for each fixed (Δ_ϕ, Λ) we assume a scalar gap Δ_0, export the resulting polynomial matrix program, convert it with `pmp2sdp`, solve feasibility with SDPB, and then use a binary search in Δ_0 to locate the critical bound. Repeating this over a scan of Δ_ϕ gives the bound curve $\Delta_0^{\text{crit}}(\Delta_\phi)$. Figure 2 shows our completed SDPB curves for $\Lambda = 6, 8, 10$. The scan used 19 values of Δ_ϕ, sampled coarsely away from the Ising point and more finely near $\Delta_\phi = \frac{1}{8}$, and a gap-search tolerance of order 10^{-2}. On our machine, the three completed runs took approximately 3 minutes, 72 minutes, and 10.5 hours respectively. Even at this modest resolution, the kink is already visible near the 2d Ising point $(\Delta_\phi, \Delta_0) = (\frac{1}{8}, 1)$.

The accuracy of our run is limited by the computational demand of the full SDPB calculation. A more accurate and high-resolution computation is shown in Figure 27 of [4], which includes higher-Λ curves (up to $\Lambda = 28$) and a broader comparison to minimal-model data. The convergence at the 2d Ising point $\Delta_\phi = \frac{1}{8}$ is striking: the bound Δ_0^{crit} tightens from 1.020 at $\Lambda = 6$ to 1.0027 at

$\Lambda = 8$, 1.00053 at $\Lambda = 12$, 1.000043 at $\Lambda = 16$, and ~ 1.0000005 at $\Lambda = 28$—within 5×10^{-7} of the exact value $\Delta_\epsilon = 1$. The bound thus appears to be saturated by the 2d Ising CFT. Our $\Lambda = 10$ run reproduces this qualitative kink structure but at a significantly looser quantitative bound; the rapid (apparently exponential) convergence with Λ is itself one of the most striking features of the numerical bootstrap.

The same procedure works in $d = 3$, where conformal blocks must be computed numerically via the recursion relations and radial expansion discussed above. In this case the bootstrap yields an upper bound on the dimension Δ_ϵ of the lowest-dimension scalar in the $\sigma \times \sigma$ OPE as a function of Δ_σ, and the resulting bound again has a sharp kink—now at $\Delta_\sigma \approx 0.518$. The location of the kink corresponds precisely to the critical 3d Ising model, in which σ is the spin operator (the $\mathbb{Z}_2$-odd scalar) and ϵ is the energy operator (the lowest $\mathbb{Z}_2$-even scalar in the $\sigma \times \sigma$ OPE).

Again, the kink is significant because the 3d Ising CFT appears to *exactly saturate* the bootstrap bound—it lives at the boundary of the allowed region in the $(\Delta_\sigma, \Delta_\epsilon)$ plane. By zooming into the kink and combining with additional constraints (such as imposing a gap in the spin-2 sector above the stress tensor, or imposing the $\mathbb{Z}_2$ symmetry that distinguishes the $\sigma \times \sigma$ and $\sigma \times \epsilon$ OPEs), the bootstrap has determined the critical exponents of the 3d Ising model to extraordinary precision [10]:

$$\Delta_\sigma = 0.5181489(10), \qquad \Delta_\epsilon = 1.412625(10). \tag{45}$$

These correspond to the traditional critical exponents via $\eta = 2\Delta_\sigma - 1$ and $\nu = 1/(3 - \Delta_\epsilon)$, giving $\eta = 0.0362978(20)$ and $\nu = 0.629971(4)$. These are the most precise determinations of 3d Ising critical exponents by any method, surpassing decades of Monte Carlo and perturbative RG results. The conformal bootstrap thus provides a concrete demonstration of how crossing symmetry and unitarity, combined with no Lagrangian input whatsoever, can pin down the properties of a strongly coupled theory with remarkable precision.

3.7 Extensions and Generalizations

The bootstrap program extends naturally beyond the case of identical scalars treated above, though the technical complexity increases. For *non-identical external operators*, the crossing equation involves multiple four-point functions (e.g., $\langle\sigma\sigma\sigma\sigma\rangle$, $\langle\sigma\sigma\epsilon\epsilon\rangle$, $\langle\epsilon\epsilon\epsilon\epsilon\rangle$ in a theory with $\mathbb{Z}_2$ symmetry) that must be simultaneously consistent. The OPE coefficients are no longer squared—they appear as products $f_{12k}\, f_{34k}$—so the positivity argument is more subtle, typically requiring semidefinite (rather than linear) constraints. This "mixed correlator" approach was essential to obtaining the *precision* Ising results quoted above. The single-correlator bootstrap of $\langle\sigma\sigma\sigma\sigma\rangle$ alone already produces the celebrated kink at $\Delta_\sigma \approx 0.518$, but it yields only a one-dimensional bound, not a small island. Isolating the Ising CFT to a tiny region in the $(\Delta_\sigma, \Delta_\epsilon)$ plane—and extracting critical exponents to six or more digits—requires this mixed-correlator system.

For external operators with spin, the conformal blocks become matrix-valued objects corresponding to the different tensor structures that can appear in three-point functions of spinning operators. The number of independent structures grows with spin, making the computation substantially more involved. Nonetheless, the bootstrap has been applied to systems with conserved currents and stress tensors as external operators, yielding bounds on central charges and current central charges [4].

The bootstrap has been applied to a wide range of physical systems beyond the 3d Ising model. A natural class of targets are the $O(N)$ vector models [9], which describe the fixed points of systems

with N-component order parameters. The bootstrap provides an alternative to the large-N saddle-point approach, working at finite N directly with no expansion parameter. For the $O(2)$ model (relevant to the superfluid helium transition) and the $O(3)$ model (the Heisenberg ferromagnet), where the $1/N$ expansion is unreliable, the mixed-correlator bootstrap has determined the leading critical exponents to a high precision.

Beyond the $O(N)$ models, the bootstrap has been applied to theories with fermions and gauge fields. In higher dimensions, it constrains 4d $\mathcal{N} = 1, 2, 4$ superconformal field theories, where superconformal symmetry provides additional structure that sharpens the bounds (see [4] for a review). It has also been extended to theories with boundaries and defects [12], and to the study of large-spin operators via the analytic bootstrap [11].

4 Future Directions

The conformal bootstrap remains a rapidly developing field with many open questions. We highlight several directions of current and future research.

Uniqueness of the 3d Ising CFT. The numerical bootstrap has isolated the 3d Ising CFT to a tiny island in the $(\Delta_\sigma, \Delta_\epsilon)$ plane, and the island continues to shrink as computational power increases—recent work [5] has pushed the precision to $\Delta_\sigma = 0.518148806(24)$ and $\Delta_\epsilon = 1.41262528(29)$. However, it remains an open question whether the bootstrap bounds converge to a *point* in the limit $\Lambda \to \infty$, which would constitute a nonperturbative proof of uniqueness for the 3d Ising CFT as the only choice of CFT data consistent with $d = 3$, $\mathbb{Z}_2$ global symmetry, and the existence of a single relevant $\mathbb{Z}_2$-odd scalar σ and a single relevant $\mathbb{Z}_2$-even scalar ϵ. This would be a clear demonstration of universality from the bootstrap axioms alone.

Analytic bootstrap and large-spin expansions. Complementary to the numerical program, the *analytic bootstrap* exploits the structure of crossing symmetry in various limits. For example, the lightcone limit ($z \to 0$ with $\bar{z}$ fixed) probes the large-spin spectrum [11] and has led to the development of large-spin perturbation theory, which determines the anomalous dimensions and OPE coefficients of operators at large spin as an asymptotic expansion.

Gauge theories and non-Lagrangian CFTs. Extending the bootstrap to theories with gauge symmetry remains a major challenge. Recent progress on bootstrapping correlators of non-abelian conserved currents in $d = 3$ [14, 15] suggests that bootstrap methods may eventually constrain strongly coupled gauge theories such as 3d QED and QCD_3, and potentially provide rigorous assessments of confinement scenarios. The bootstrap is also uniquely suited to studying non-Lagrangian CFTs—theories that have no known weakly coupled description—since it requires no Lagrangian input.

Defects, boundaries, and non-trivial geometries. The bootstrap has been extended to CFTs with boundaries [12], line defects, and surface operators, each of which introduces additional crossing equations involving "bulk-to-defect" OPE coefficients. These setups are physically relevant: boundary CFTs describe surface critical phenomena, and line defects model impurities in critical systems. The defect bootstrap is still in its early stages relative to the bulk program.

Computational methods. On the algorithmic side, the development of more efficient SDP solvers, better strategies for choosing which correlators and derivative orders to study, and the use of machine learning to navigate the space of bootstrap bounds are all active areas. The optimal choice of functional basis—whether derivatives at the crossing-symmetric point, analytic functionals, or other bases—remains an open problem.

Applications to AdS/CFT. Chandra, Collier, Hartman, and Maloney [13] have defined an ensemble of large-c CFT data by averaging over OPE coefficients subject to bootstrap constraints, and shown that calculations in this ensemble reproduce semiclassical 3d gravity. A key open question is whether imposing the full set of bootstrap constraints singles out individual CFTs from the ensemble, which would amount to a nonperturbative definition of the holographic dual.

5 Conclusions

We have developed the formalism of conformal field theory in $d \geq 3$ dimensions and derived the conformal bootstrap. Starting from the conformal algebra $SO(d+1,1)$ and its representation theory, we showed how conformal symmetry fixes two- and three-point functions up to the CFT data $\{\Delta_i, f_{ijk}\}$, and constrains four-point functions to depend on an arbitrary function of two cross ratios. Radial quantization and the state-operator correspondence provided the Hilbert space structure needed to establish unitarity bounds and the completeness of the primary-descendant decomposition. From these we derived the operator product expansion which, combined with the orthonormality guaranteed by reflection positivity, reduced all correlation functions to the CFT data outlined above.

The conformal bootstrap is a requirement on the associativity of the OPE: different channels of the four-point function must agree, leading to the crossing equation. The positivity of OPE coefficients squared, a consequence of unitarity, converts this into a geometric constraint on a function space that can be attacked by searching for linear functionals. We worked through this argument explicitly in two stages. First, we constructed a low-dimensional geometric bound in $d = 2$ by projecting crossing vectors to $\mathbb{R}^2$ at fixed $\Delta_\phi = 0.20$. Second, we implemented the SDPB-based Section 10.7 workflow: for each (Δ_ϕ, Λ) we assumed a scalar gap Δ_0, converted the resulting polynomial matrix program with `pmp2sdp`, solved feasibility with `SDPB`, and used binary search in Δ_0 to obtain bound curves $\Delta_0^{\text{crit}}(\Delta_\phi)$. This produces the same qualitative kink structure near the 2d Ising point as the benchmark analysis.

The connection to the 3d case is structural rather than accidental: the same pipeline (crossing equations, truncated derivative functionals, SDP feasibility, and scans over external/operator dimensions) underlies the 3d Ising bootstrap. In 2d, the kink near $(\Delta_\phi, \Delta_0) = (\frac{1}{8}, 1)$ is a controlled testbed where exact CFT data are known, while in 3d the corresponding kink near $\Delta_\sigma \approx 0.518$ identifies the physical Ising CFT and, with mixed-correlator constraints, sharpens into a small island. Thus the 2d and 3d computations are best viewed as the same numerical method applied at different levels of complexity: in 2d as validation and intuition, and in 3d as a precision tool for extracting critical exponents.

The conformal bootstrap demonstrates that the basic mathematical ingredients of CFT—unitarity, crossing symmetry, and the operator product expansion—are far more constraining than one might naively expect. Without writing down a Lagrangian, performing a path integral, or summing Feynman diagrams, one can extract precise quantitative predictions about strongly

coupled systems. This is a profound illustration of the power of consistency conditions in quantum field theory, and the many open directions discussed above suggest that the bootstrap program will continue to produce important results for years to come.

A Additional Projection Plots

This appendix collects several additional two-dimensional projections of the crossing vectors for identical scalar fields. These illustrate how the qualitative geometry changes as the projection map is varied.

B AI Statement

AI was used for transcribing some of the longer equations and finding/parsing references. It was particularly helpful for some of the future directions and related papers to read about. AI also edited and suggested changes for clarity/spelling/grammar, and it wrote/debugged a lot of the code that went into generating those figures, though I picked the parameters and played around with different settings to obtain good figures.

References

[1] S. Ferrara, A. F. Grillo, and R. Gatto, "Tensor representations of conformal algebra and conformally covariant operator product expansion," *Ann. Phys.* **76**, 161 (1973).

[2] A. M. Polyakov, "Nonhamiltonian approach to conformal quantum field theory," *Zh. Eksp. Teor. Fiz.* **66**, 23 (1974).

[3] R. Rattazzi, V. S. Rychkov, E. Tonni, and A. Vichi, "Bounding scalar operator dimensions in 4D CFT," *JHEP* **0812**, 031 (2008), arXiv:0807.0004.

[4] D. Simmons-Duffin, "TASI Lectures on the Conformal Bootstrap," arXiv:1602.07982.

[5] N. Su, M. Go, and Y. He, "Bootstrapping the 3d Ising stress tensor," *JHEP* **2025**, 136 (2025), arXiv:2411.15300.

[6] S. El-Showk and M. F. Paulos, "Bootstrapping conformal field theories with the extremal functional method," *Phys. Rev. Lett.* **111**, 241601 (2013), arXiv:1211.2810.

[7] D. Simmons-Duffin, "A semidefinite program solver for the conformal bootstrap," *JHEP* **1506**, 174 (2015), arXiv:1502.02033.

[8] A. B. Zamolodchikov, "Conformal symmetry in two-dimensional space: recursion representation of conformal block," *Theor. Math. Phys.* **73**, 1088 (1987).

[9] F. Kos, D. Poland, and D. Simmons-Duffin, "Bootstrapping the $O(N)$ vector models," *JHEP* **1406**, 091 (2014), arXiv:1307.6856.

[10] F. Kos, D. Poland, D. Simmons-Duffin, and A. Vichi, "Precision islands in the Ising and $O(N)$ models," *JHEP* **1608**, 036 (2016), arXiv:1603.04436.

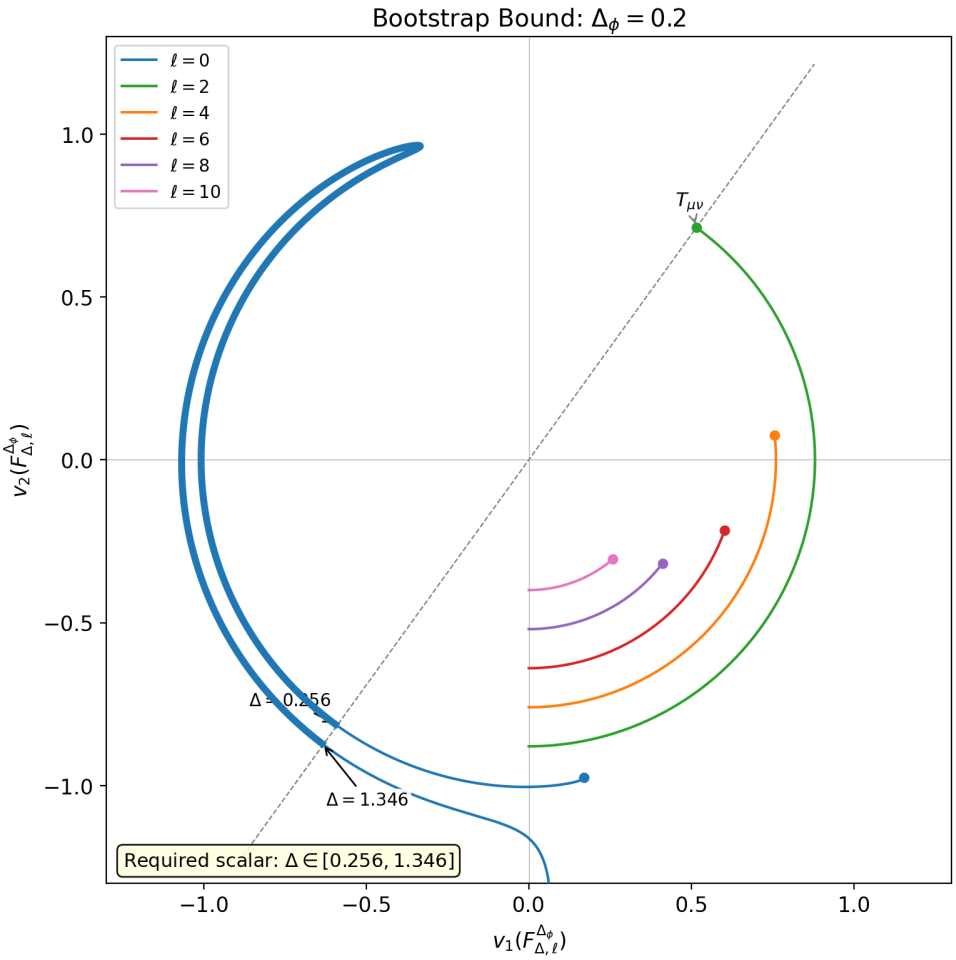

Figure 3: Alternative projection with $\Delta_\phi = 0.20$ and $\vec{v}(F) = \big(H(0.5, 0.70) - H(0.5, 0.35),\ H(0.40, 0.65) - H(0.26, 0.30)\big)$.

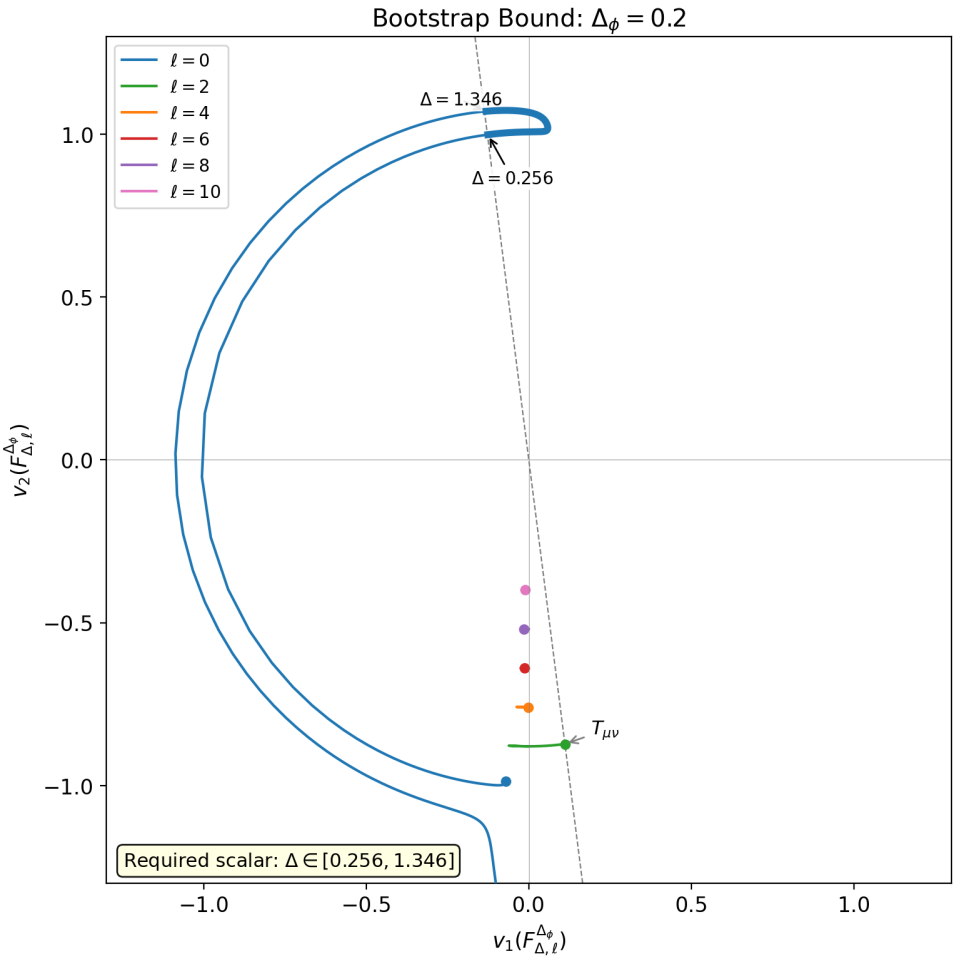

Figure 4: Alternative projection with $\Delta_\phi = 0.20$ and $\vec{v}(F) = \big(H(0.48, 0.66) - H(0.48, 0.32),\ H(0.42, 0.57) - H(0.31, 0.22)\big)$.

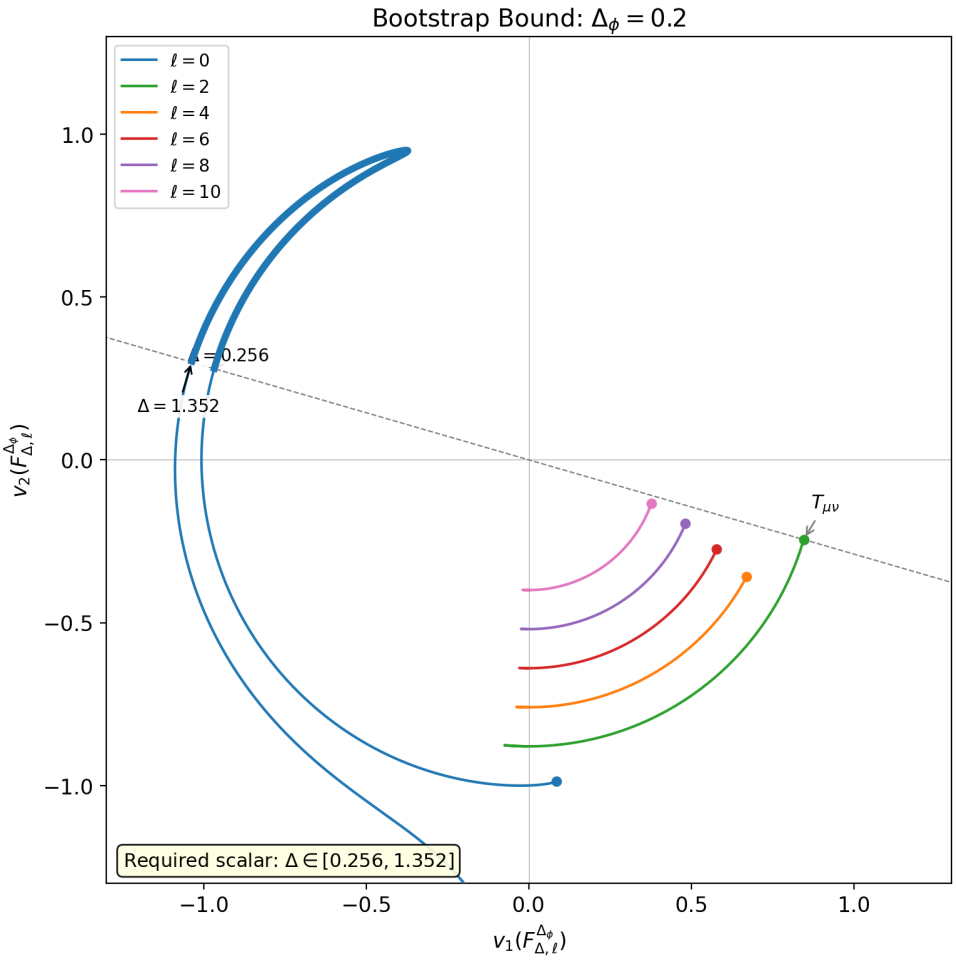

Figure 5: Alternative projection with $\Delta_\phi = 0.20$ and $\vec{v}(F) = \big(H(0.7755, 0.2781) - H(0.7799, 0.4173),\ H(0.4840, 0.3756) - H(0.7623, 0.7749)\big)$.

[11] Z. Komargodski and A. Zhiboedov, "Convexity and liberation at large spin," *JHEP* **1311**, 140 (2013), arXiv:1212.4103.

[12] P. Liendo, L. Rastelli, and B. C. van Rees, "The bootstrap program for boundary CFT$_d$," *JHEP* **1307**, 113 (2013), arXiv:1210.4258.

[13] J. Chandra, S. Collier, T. Hartman, and A. Maloney, "Semiclassical 3D gravity as an average of large-c CFTs," *JHEP* **2212**, 069 (2022), arXiv:2203.06511.

[14] S. M. Chester and S. S. Pufu, "Towards bootstrapping QED$_3$," *JHEP* **1608**, 019 (2016), arXiv:1601.03476.

[15] A. Atanasov, A. Hillman, D. Poland, J. Rong, and N. Su, "Non-abelian currents bootstrap," *JHEP* **2403**, 175 (2024), arXiv:2302.11585.