

Solving the 3d Ising Model with the Conformal Bootstrap

II. c -Minimization and Precise Critical Exponents

Sheer El-Showk^a, Miguel F. Paulos^b, David Poland^c,

Slava Rychkov^{a,d}, David Simmons-Duffin^e, Alessandro Vichi^f

^a *Theory Division, CERN, Geneva, Switzerland*

^b *Department of Physics, Brown University, Box 1843, Providence, RI 02912-1843, USA*

^c *Department of Physics, Yale University, New Haven, CT 06520, USA*

^d *Faculté de Physique, Université Pierre et Marie Curie
& Laboratoire de Physique Théorique, École Normale Supérieure, Paris, France*

^e *School of Natural Sciences, Institute for Advanced Study,
Princeton, New Jersey 08540, USA*

^f *Theoretical Physics Group, Ernest Orlando Lawrence Berkeley National Laboratory
& Center for Theoretical Physics, University of California, Berkeley, CA 94720, USA*

Abstract

We use the conformal bootstrap to perform a precision study of the operator spectrum of the critical 3d Ising model. We conjecture that the 3d Ising spectrum minimizes the central charge c in the space of unitary solutions to crossing symmetry. Because extremal solutions to crossing symmetry are uniquely determined, we are able to precisely reconstruct the first several \mathbb{Z}_2 -even operator dimensions and their OPE coefficients. We observe that a sharp transition in the operator spectrum occurs at the 3d Ising dimension $\Delta_\sigma = 0.518154(15)$, and find strong numerical evidence that operators decouple from the spectrum as one approaches the 3d Ising point. We compare this behavior to the analogous situation in 2d, where the disappearance of operators can be understood in terms of degenerate Virasoro representations.

March 2014

To appear in a special issue of J.Stat.Phys. in memory of Kenneth Wilson

Contents

1	Introduction	2
2	A Conjecture for the 3d Ising Spectrum	3
2.1	Brief CFT Reminder	5
2.2	The Space of Unitary, Crossing-Symmetric 4-point Functions	5
2.3	Approaching the Boundary of $\mathcal{C}_{\Delta_\sigma}$	7
2.4	Equivalence of Δ_ϵ -maximization and c -minimization	8
3	The 3d Ising Spectrum from c-minimization	12
3.1	Δ_σ and c	12
3.2	Extracting the Spectrum and OPE Coefficients in $\sigma \times \sigma$	13
3.3	The Leading \mathbb{Z}_2 -Even Scalar ϵ	14
3.4	Higher Scalars: General Features	16
3.5	The Second \mathbb{Z}_2 -Even Scalar ϵ'	19
3.6	Spin 2 Operators	19
3.7	Higher spins	21
4	Comparison to Results by Other Techniques	23
5	2d Checks	26
5.1	Spectrum Extraction at Fixed $\Delta_\sigma = 1/8$	26
5.2	c -Minimization	28
5.3	Interpolating Between Minimal Models	32
6	Technical Details	34
6.1	Partial Fraction Representation for Conformal Blocks	34
6.1.1	The Casimir Equation	36
6.1.2	Recursion Relations	37
6.2	A Customized Simplex Method	38
6.2.1	The Primal Simplex Method	39
6.2.2	Choosing an Initial Vertex	40

6.2.3	Adaptation to Optimization over $\mathcal{C}_{\Delta_\sigma}$	41
6.2.4	A Strategy for Reduced Cost Minimization	43
6.2.5	Hot Start from the Mean Field Theory Spectrum	45
6.2.6	Linear Functionals	45
6.2.7	Δ_ϵ -Maximization	47
6.3	Implementation and Running Time Details	48
7	Summary and Discussion	48
A	Comparison to Semidefinite Programming	51

1 Introduction

In [1], we initiated the conformal bootstrap approach to studying the “3d Ising CFT”—the CFT describing the 3d Ising model at criticality. This CFT seems to occupy a very special place in the space of unitary \mathbb{Z}_2 -symmetric 3d CFTs. In the plane parametrized by the leading \mathbb{Z}_2 -odd and \mathbb{Z}_2 -even scalar dimensions Δ_σ and Δ_ϵ , it seems to live at a “kink” on the boundary of the region allowed by the constraints of crossing symmetry and unitarity. Turning this around, the position of this kink can be used to determine Δ_σ and Δ_ϵ , giving results in agreement with the renormalization group and lattice methods. Moreover, compared to these other techniques our method has two important advantages. Firstly, the needed computational time scales favorably with the final accuracy (the kink localization), promising high-precision results for the critical exponents. Secondly, due to the fact that the many parameters characterizing the CFT talk to each other in the bootstrap equation, the results show interesting structure and interrelations. The very existence of the kink tying Δ_σ to Δ_ϵ is one such interrelation, and there are many more [1]. For example, previously we found that the subleading \mathbb{Z}_2 -even scalar and spin 2 operator dimensions show interesting variation near the kink, while the stress tensor central charge c seems to take a minimal value there.

Exploiting these advantages of the bootstrap approach will be a major theme of this second work in the series. Our primary goal will be to use the conformal bootstrap to make a high precision determination of all low-lying operator dimensions and OPE coefficients in the 3d Ising CFT. In doing so, we would also like to gain some insights into why the 3d Ising solution is special. We will make extensive use of the fact that solutions living on the boundary of the region allowed by crossing symmetry and unitarity can be uniquely reconstructed [2]. In order to reach the boundary, we will work with the conjecture that the central charge c is minimized for the 3d Ising CFT. This conjecture combined with uniqueness of the boundary solution allows us to numerically reconstruct the solution to crossing symmetry for different values of Δ_σ , using a modification of Dantzig’s simplex algorithm.

We find that the resulting \mathbb{Z}_2 -even spectrum shows a dramatic transition in the vicinity of $\Delta_\sigma = 0.518154(15)$, giving a high precision determination of the leading critical exponent η . Focusing on the transition region, we are able to extract precise values of the first several \mathbb{Z}_2 -even operator dimensions and of their OPE coefficients, see Table 1. We also give reasonable estimates for the locations of all low dimension ($\Delta \lesssim 13$) scalar and spin 2 operators in the \mathbb{Z}_2 -even spectrum.

spin & \mathbb{Z}_2	name	Δ	OPE coefficient
$\ell = 0, \mathbb{Z}_2 = -$	σ	0.518154(15)	
$\ell = 0, \mathbb{Z}_2 = +$	ϵ	1.41267(13)	$f_{\sigma\sigma\epsilon}^2 = 1.10636(9)$
	ϵ'	3.8303(18)	$f_{\sigma\sigma\epsilon'}^2 = 0.002810(6)$
$\ell = 2, \mathbb{Z}_2 = +$	T	3	$c/c_{\text{free}} = 0.946534(11)$
	T'	5.500(15)	$f_{\sigma\sigma T'}^2 = 2.97(2) \times 10^{-4}$

Table 1: Precision information about the low-lying 3d Ising CFT spectrum and OPE coefficients extracted in this work. See sections 3.4 and 3.6 for preliminary information about higher-dimension $\ell = 0$ and $\ell = 2$ operators. See also section 4 for a comparison to results by other techniques.

The transition also shows the highly intriguing feature that certain operators *disappear* from the spectrum as one approaches the 3d Ising point. This decoupling of states gives an important characterization of the 3d Ising CFT. This is similar to what occurs in the 2d Ising model, where the decoupling of operators can be rigorously understood in terms of degenerate representations of the Virasoro symmetry. To better understand this connection, we give a detailed comparison to the application of our c -minimization algorithm in 2d, where the exact spectrum of the 2d Ising CFT and its interpolation through the minimal models is known. We conclude with a discussion of important directions for future research.

2 A Conjecture for the 3d Ising Spectrum

Consider a 3d CFT with a scalar primary operator σ of dimension Δ_σ . In [1], we studied the constraints of crossing symmetry and unitarity on the four-point function $\langle\sigma\sigma\sigma\sigma\rangle$. From these constraints, we derived universal bounds on dimensions and OPE coefficients of operators appearing in the $\sigma \times \sigma$ OPE. Figure 1, for example, shows an upper bound on the dimension of the lowest-dimension scalar in $\sigma \times \sigma$ (which we call ϵ), as a function of Δ_σ . This bound is a consequence of very general principles - conformal invariance, unitarity, and crossing symmetry - yet it has a striking “kink” near $(\Delta_\sigma, \Delta_\epsilon) \approx (0.518, 1.412)$, indicating that these dimensions have special significance in the space of 3d CFTs. Indeed, they are believed to be realized in the 3d Ising CFT.

The curves in Figure 1 are part of a family of bounds labeled by an integer N (defined in section 2.3), which get stronger as N increases. It appears likely that the 3d Ising CFT

saturates the optimal bound on Δ_ϵ , achieved in the limit $N \rightarrow \infty$. Thus, in this work, we will take seriously the idea:

- $\langle \sigma\sigma\sigma\sigma \rangle$ in the 3d Ising CFT lies on the boundary of the space of unitary, crossing symmetric four-point functions.

We will further present a plausible guess for where on the boundary the 3d Ising CFT lies, and use this to formulate a precise conjecture for the spectrum of operators in the $\sigma \times \sigma$ OPE.




Figure 1: An upper bound on the dimension of the lowest dimension scalar $\epsilon \in \sigma \times \sigma$, as a function of Δ_σ . The blue shaded region is allowed; the white region is disallowed. *Left:* The bound at $N = 78$ [1]. *Right:* The bound at $N = 105$, in the region near the kink. This bound is thus somewhat stronger than the previous one, and the kink is sharper.

Although the 3d Ising CFT is certainly special, it is perhaps surprising that one might find it by studying a single four-point function. After all, the full consistency constraints of a CFT include crossing symmetry and unitarity for *every* four-point function, including all possible operators in the theory. Nevertheless, other recent work supports the idea that for some special CFTs it may be enough to consider $\langle \sigma\sigma\sigma\sigma \rangle$. For example, one can compute similar bounds in fractional spacetime dimension $2 \leq d \leq 4$. These bounds have similar kinks which agree with the operator dimensions present at the Wilson-Fisher fixed point near $d = 4$ and the 2d Ising CFT when $d = 2$ [3]. An analogous story holds for theories with $O(n)$ global symmetry in 3d, where $O(n)$ vector models appear to saturate their associated dimension bounds [4].

As a check on our conjecture, we will also apply it to the 2d Ising CFT. We find good agreement with the known exact solution, and previous numerical explorations of the 2d bootstrap [2]. Our study of 2d will serve as a useful guide for interpreting our results in 3d.

2.1 Brief CFT Reminder

Let us first recall some basic facts about CFT four-point functions that we will need in our analysis. A four-point function of a real scalar primary σ in a CFT takes the form

$$\langle \sigma(x_1)\sigma(x_2)\sigma(x_3)\sigma(x_4) \rangle = \frac{1}{x_{12}^{2\Delta_\sigma} x_{34}^{2\Delta_\sigma}} g(u, v), \quad (2.1)$$

where $u = \frac{x_{12}^2 x_{34}^2}{x_{13}^2 x_{24}^2}$ and $v = \frac{x_{14}^2 x_{23}^2}{x_{13}^2 x_{24}^2}$ are conformally invariant cross-ratios. The function $g(u, v)$ can be expanded in conformal blocks

$$g(u, v) = \sum_{\mathcal{O} \in \sigma \times \sigma} p_{\Delta, \ell} G_{\Delta, \ell}(u, v), \quad (2.2)$$

where the sum is over primary operators \mathcal{O} with dimension Δ and spin ℓ appearing in the $\sigma \times \sigma$ OPE, and the coefficients $p_{\Delta, \ell} = f_{\sigma\sigma\mathcal{O}}^2$ are positive. Only even spins ℓ can appear in $\sigma \times \sigma$, and Δ must obey the unitarity bounds

$$\Delta \geq \begin{cases} (d-2)/2 & \text{if } \ell = 0 \text{ (excluding the unit operator),} \\ \ell + d - 2 & \text{if } \ell > 0, \end{cases} \quad (2.3)$$

where d is the spacetime dimension (we will mostly be interested in $d = 3$). We normalize σ so that the OPE coefficient of the unit operator is $f_{\sigma\sigma 1} = 1$.

Finally, invariance of the four-point function (2.1) under permutations of the x_i implies the crossing-symmetry constraint

$$g(u, v) = \left(\frac{u}{v}\right)^{\Delta_\sigma} g(v, u). \quad (2.4)$$

All results in this paper will be based on Eqs. (2.2), (2.4) and on the information about conformal blocks reviewed in section 6.1.

2.2 The Space of Unitary, Crossing-Symmetric 4-point Functions

Instead of focusing on a specific CFT, let us turn these facts around and consider all possible scalar four-point functions in any CFT. Let $\mathcal{C}_{\Delta_\sigma}$ be the space of linear combinations of conformal blocks

$$g(u, v) = \sum_{\Delta, \ell} p_{\Delta, \ell} G_{\Delta, \ell}(u, v) \quad (2.5)$$

such that

1. (Δ, ℓ) obey the unitarity bounds (2.3),
2. $p_{0,0} = 1$,
3. $p_{\Delta, \ell} \geq 0$,

4. $g(u, v)$ is crossing-symmetric, Eq. (2.4).

We include the second condition because the unit operator should be present with $f_{\sigma\sigma 1} = 1$. The third condition is because the OPE coefficients $f_{\sigma\sigma\mathcal{O}}$ are real in a unitary theory, so their squares should be positive. Note that $\mathcal{C}_{\Delta_\sigma}$ depends on the parameter Δ_σ through the crossing constraint (2.4).¹

Let us make a few comments about this space. Firstly, $\mathcal{C}_{\Delta_\sigma}$ is convex. This follows because positive linear combinations $tg_1(u, v) + (1-t)g_2(u, v)$, of four-point functions g_1, g_2 also satisfy the above conditions. Geometrically, we can think of $\mathcal{C}_{\Delta_\sigma}$ as a convex cone given by the unitarity condition $p_{\Delta,\ell} \geq 0$, subject to the affine constraints $p_{0,0} = 1$ and Eq. (2.4). This picture is illustrated in Figure 2.




Figure 2: The space $\mathcal{C}_{\Delta_\sigma}$ (in blue) is the intersection of the convex cone given by the unitarity conditions $p_{\Delta,\ell} \geq 0$ with a hyperplane given by the affine constraints $p_{0,0} = 1$ and Eq. (2.4). It always contains a point corresponding to Mean Field Theory, and might contain points corresponding to other CFTs with a scalar of dimension Δ_σ . We conjecture that for a special value of Δ_σ , the 3d Ising CFT lies on the boundary of $\mathcal{C}_{\Delta_\sigma}$.

Secondly, $\mathcal{C}_{\Delta_\sigma}$ is nonempty. If a CFT with a scalar operator of dimension Δ_σ exists, then $\langle\sigma\sigma\sigma\sigma\rangle$ in that theory certainly gives a point in $\mathcal{C}_{\Delta_\sigma}$. Furthermore, $\mathcal{C}_{\Delta_\sigma}$ always contains a point corresponding to Mean Field Theory (MFT), where the four-point function $\langle\sigma\sigma\sigma\sigma\rangle$ is given by a sum of products of two-point functions.² The MFT four-point function is crossing-symmetric and can be expanded as a sum of conformal blocks with positive coefficients. The explicit coefficients $p_{\Delta,\ell}^{\text{MFT}}$ can be found in [6, 7].³

¹These conditions are the obvious ones implied by conformal symmetry, unitarity, and crossing symmetry of $\langle\sigma\sigma\sigma\sigma\rangle$. We expect that $p_{\Delta,\ell}$'s in actual CFTs satisfy further conditions related to consistency of other four-point functions, and perhaps more exotic conditions like consistency of the theory on compact manifolds. However, we do not impose these constraints in defining $\mathcal{C}_{\Delta_\sigma}$. It will be interesting and important to explore them in future work, for instance [5].

²Such a scalar σ is sometimes called a generalized free field.

³MFT is not a genuine local CFT because it doesn't contain a stress-tensor. However, it does appear as a point in $\mathcal{C}_{\Delta_\sigma}$.

What is the dimensionality of $\mathcal{C}_{\Delta_\sigma}$? The answer to this question is not immediately obvious, since it is an infinite dimensional cone intersected with an infinite dimensional hyperplane. For Δ_σ at the unitarity bound, $\Delta_\sigma = \frac{d-2}{2}$, it is possible to show that $\mathcal{C}_{\Delta_\sigma}$ consists of only one point: the free scalar four-point function. On the other hand, $\mathcal{C}_{\Delta_\sigma}$ is likely to be an infinite-dimensional space for all $\Delta_\sigma > \frac{d-2}{2}$. When $\Delta_\sigma > d-2$, we can demonstrate this rigorously by exhibiting an infinite set of linearly independent crossing-symmetric, unitary four-point functions. Namely, consider the tensor product of Mean Field Theories with scalars ϕ_1 and ϕ_2 of dimensions δ and $\Delta_\sigma - \delta$, and let $\sigma = \phi_1\phi_2$. This gives a family of linearly-independent four-point functions labeled by a continuous parameter $\frac{d-2}{2} \leq \delta \leq \Delta_\sigma/2$. While this argument does not apply for $\frac{d-2}{2} < \Delta_\sigma \leq d-2$, we believe that $\mathcal{C}_{\Delta_\sigma}$ remains infinite-dimensional in this range. Some numerical evidence for this will be discussed in section 6.2.6.

2.3 Approaching the Boundary of $\mathcal{C}_{\Delta_\sigma}$

Every point on the boundary of a convex space is the maximum of some linear function over that space. Conversely, if we have a bounded linear function on a convex space, then the maximum of that function is well-defined and generically unique. Assuming that the 3d Ising CFT lies on the boundary of $\mathcal{C}_{\Delta_\sigma}$, we should ask: *what does the 3d Ising CFT maximize?*

Before addressing this question, let us introduce some details about optimization over $\mathcal{C}_{\Delta_\sigma}$.⁴ To explore $\mathcal{C}_{\Delta_\sigma}$ numerically, we construct it via a limiting procedure where we truncate the crossing equation (2.4) to a finite number of constraints.⁵ Specifically, define $\mathcal{C}_{\Delta_\sigma}^{(N)}$ by the same conditions as $\mathcal{C}_{\Delta_\sigma}$, except with the crossing-symmetry constraint replaced by

$$\partial_u^m \partial_v^n \left(g(u, v) - \left(\frac{u}{v} \right)^{\Delta_\sigma} g(v, u) \right) \Big|_{u=v=1/4} = 0 \quad (2.6)$$

for N different (nonzero) derivatives (m, n) .⁶ The crossing-symmetric point $u = v = 1/4$ is chosen as usual in the numerical bootstrap, to optimize the convergence of the conformal block expansion (2.5) [8, 9]. This gives a sequence of smaller and smaller convex spaces, with $\mathcal{C}_{\Delta_\sigma}$ as a limit,

$$\mathcal{C}_{\Delta_\sigma}^{(1)} \supset \mathcal{C}_{\Delta_\sigma}^{(2)} \supset \dots \supset \mathcal{C}_{\Delta_\sigma}. \quad (2.7)$$

Consider a function f over $\mathcal{C}_{\Delta_\sigma}^{(N)}$ which is maximized at a point g_N on the boundary. Generically, g_N will not lie inside $\mathcal{C}_{\Delta_\sigma}$. However, by following the sequence g_N as $N \rightarrow \infty$, we can approach the maximum g_* over $\mathcal{C}_{\Delta_\sigma}$ (Figure 3). It's important that for functions considered below, the optimal point will be uniquely determined, i.e. the maximum is not

⁴We give a full description of our algorithm in section 6.2.

⁵On the other hand, we will *not* have to restrict the dimensions Δ to a discrete set, as was done originally in [8]. For this work, we have implemented an optimization algorithm (described in section 6.2) that works with continuously varying Δ .

⁶This equation is schematic. In section 6.2.3 we will pass from u, v to the variables z, \bar{z} , and take N partial derivatives with respect to those variables.

degenerate. This will allow us to pick out, for each Δ_σ and N , the optimal spectrum and the OPE coefficients. We will see in section 6.2.6 that the optimal spectrum at each N contains order N operators. As N increases, we observe rapid convergence of lower dimension operators and also growth of the number of operators present. The hope is that the low lying 3d Ising CFT spectrum can be recovered in the limit. Here we are taking inspiration and encouragement from [2], where it was shown that in 2d a similar procedure recovers the exactly known spectrum of the 2d Ising CFT.




Figure 3: Each number of derivatives N gives an approximation to $\mathcal{C}_{\Delta_\sigma}$ that shrinks as N increases. By repeatedly solving an optimization problem for each N , we can follow a path to the boundary of $\mathcal{C}_{\Delta_\sigma}$.

Note that while g_N doesn't necessarily lie inside $\mathcal{C}_{\Delta_\sigma}$, the value $f(g_N)$ is a rigorous upper bound on $\max_{\mathcal{C}_{\Delta_\sigma}}(f)$. Further, this bound improves as $N \rightarrow \infty$. The bounds $f(g_N)$ are precisely those usually discussed in the conformal bootstrap literature (and shown for example in Figures 1 and 4). When f is linear, the procedure of maximizing over $\mathcal{C}_{\Delta_\sigma}$ is related by linear programming duality to the procedure of minimizing linear functionals, going back to [8]. From g_N , one can always construct an “extremal” linear functional touching the boundary of $\mathcal{C}_{\Delta_\sigma}^{(N)}$ at g_N , and vice versa. In [2], zeroes of the extremal functional were used to extract the spectrum of the optimal solution; knowing the spectrum, the squared OPE coefficients $p_{\Delta,\ell}$ could be determined by solving the crossing equation. In this work we will approach g_N iteratively from inside of $\mathcal{C}_{\Delta_\sigma}^{(N)}$, having a solution to crossing at each step of the maximization procedure. We will refer to the approach described here as the “direct” or “primal” method, as opposed to the “dual” method of previous studies. In principle, they are equivalent.

2.4 Equivalence of Δ_ϵ -maximization and c -minimization

Let us return to the question: *what does the 3d Ising CFT maximize?* As we have seen, there is good evidence that the 3d Ising CFT belongs to a family of solutions to crossing which, for a fixed Δ_σ , maximize Δ_ϵ (the dimension of the lowest-dimension \mathbb{Z}_2 -even scalar). However, Δ_ϵ is not a linear function on $\mathcal{C}_{\Delta_\sigma}$ (the lowest dimension in a sum of four-point functions is the minimum dimension in either four-point function). For conceptual and practical reasons, it will be useful to find an alternative characterization of the 3d Ising point in terms of a linear function on $\mathcal{C}_{\Delta_\sigma}$.

In [1], we also computed bounds on the squared OPE coefficient of the spin 2, dimension d operator. It is natural to assume that the 3d Ising CFT contains one and only one operator with such quantum numbers, namely the conserved stress tensor of the theory.⁷ The OPE coefficient $p_T \equiv p_{d,2}$ is related via Ward identities to the central charge c ,

$$c \propto \frac{\Delta_\sigma^2}{p_T}, \quad (2.8)$$

with a d -dependent factor which depends on the normalization of c .⁸ We use the definition of the central charge c as the coefficient in the two point correlation function of the stress tensor operator. This definition works for any d , in particular for $d = 3$. In $d = 2$, the central charge can also be defined as the coefficient of a central term in a Lie algebra extending the global conformal algebra (the Virasoro algebra), but an analogous interpretation for other d is unknown and likely impossible. To avoid normalization ambiguities, we will give results below for the ratio c/c_{free} , the latter being the central charge of the free scalar theory in the same number of dimensions.

Figure 4 shows this lower bound on c (equivalently, an upper bound on p_T) as a function of Δ_σ . The bound displays a sharp minimum at the same value of Δ_σ as the kink in Figure 1, suggesting that the 3d Ising CFT might also minimize c .⁹ We will call this method of localizing the 3d Ising point “ c -minimization,” although it should be kept in mind that in practice we minimize c by maximizing p_T , which is a linear function on $\mathcal{C}_{\Delta_\sigma}$.




Figure 4: A lower bound on c (the coefficient of the two-point function of the stress tensor), as a function of Δ_σ . Left: the bound from [1] computed with $N = 78$. Right: a slightly stronger bound at $N = 105$ in the region near the minimum.

We have done numerical studies of both Δ_ϵ -maximization and c -minimization, and found evidence that for all Δ_σ in the neighborhood of the expected 3d Ising value ≈ 0.518 these

⁷The 3d Ising CFT can be obtained as an IR fixed point of the ϕ^4 theory. The UV stress tensor then naturally gives rise to the IR stress tensor, but there is no reason to expect that a second operator with the same quantum numbers will emerge in the IR. This definitely does not happen in the ϵ -expansion.

⁸In our previous work [1, 3] c was denoted C_T .

⁹By contrast, bounds on c in 4d do not show a similar minimum [10–12].

optimization problems are solved at the same (uniquely-determined) point on the boundary of $\mathcal{C}_{\Delta_\sigma}$. In other words, they are equivalent conditions. As an example, one can compute the maximum possible dimension Δ_ϵ^{\max} , and also the value Δ_ϵ^* in the four-point function that minimizes c . As already mentioned, for each Δ_σ this four-point function turns out to be unique, so that Δ_ϵ^* is uniquely determined. It is then interesting to consider the difference (non-negative by definition)

$$\delta = \Delta_\epsilon^{\max} - \Delta_\epsilon^*, \quad \Delta_\epsilon^* \equiv \Delta_\epsilon|_{c \rightarrow \min}. \quad (2.9)$$

We plot this difference in Figure 5, where we have chosen $N = 105$. The difference is tiny for all values of Δ_σ , and near one point it drops sharply towards zero — as we will see below, that’s where the 3d Ising CFT sits. We expect that as one increases the precision by taking $N \rightarrow \infty$, the difference δ should go exactly to zero at the 3d Ising point.




Figure 5: The difference $\Delta_\epsilon^{\max} - \Delta_\epsilon|_{c \rightarrow \min}$ as a function of Δ_σ , computed at $N = 105$.

A natural explanation of these observations is that, for a fixed Δ_σ there exists a corner point on the boundary of $\mathcal{C}_{\Delta_\sigma}$, which can be reached by maximizing any of several different functions — in particular, both Δ_ϵ and p_T are suitable (Figure 6). Although this corner is roughly present in the approximations $\mathcal{C}_{\Delta_\sigma}^{(N)}$, it should emerge sharply as $N \rightarrow \infty$.

In what follows, we will mainly be using c -minimization rather than Δ_ϵ -maximization to localize the 3d Ising CFT. One reason for this is practical: since p_T is a linear function on $\mathcal{C}_{\Delta_\sigma}$, maximizing p_T turns out to be much easier than maximizing Δ_ϵ in our numerical approach (see section 6.2.7 for an explanation).

In the above discussion we varied Δ_σ in an interval around 0.518. The true 3d Ising value of Δ_σ should correspond to the position of kinks in the given Δ_ϵ and c bounds. Clearly, *overall* Δ_ϵ -maximization is not the right criterion to pick out the kink in the $(\Delta_\sigma, \Delta_\epsilon)$ plane (Figure 1). We could of course cook up an ad hoc linear combination of Δ_σ and Δ_ϵ whose maximum would pick out the kink. An especially neat feature of the c -minimization is that no such ad hoc procedure is necessary: the kink is at the same time the overall minimum of c when Δ_σ is allowed to vary. It is worth emphasizing that the factor Δ_σ^2 in (2.8) is important for achieving this property, i.e. without this factor $1/p_T$ would have a kink but




Figure 6: Presumed geometry of $\mathcal{C}_{\Delta_\sigma}$ near the 3d Ising CFT. Its boundary must have a corner point which can be reached by either maximizing Δ_ϵ or minimizing c .

not an overall minimum at the 3d Ising point. Thus it is really the quantity c , rather than the OPE coefficient of $T_{\mu\nu}$, which seems to enjoy a special and non-trivial role in determining Δ_σ for the 3d Ising model.

This last point allows us to conjecture that the 3d Ising CFT minimizes the central charge c over the space of unitary, crossing symmetric four-point functions. More succinctly:

$$\Delta_\sigma, g(u, v) \text{ in 3d Ising} = \underset{\Delta_\sigma, g(u, v) \in \mathcal{C}_{\Delta_\sigma}}{\operatorname{argmin}} c, \quad (2.10)$$

up to two qualifications mentioned below.

There are clear conceptual advantages to phrasing our conjecture in terms of the c minimum instead of the Δ_ϵ kink. The stress tensor $T_{\mu\nu}$ is arguably more special than the scalar operator ϵ — it is a conserved current present in every local CFT. Further, the central charge c can be interpreted loosely as a measure of how “simple” a CFT is; our conjecture then implies that the 3d Ising CFT is the “simplest” 3d CFT.¹⁰ In particular, the 3d Ising CFT is as far as possible from Mean Field Theory, which has $c = \infty$.

The two qualifications concerning (2.10) are as follows. Firstly, the minimum over Δ_σ is local; it is supposed to be taken in the region sufficiently close to the unitarity bound, $0.5 \leq \Delta_\sigma \lesssim 1$. For larger Δ_σ the bound in Figure 4 curves down and approaches zero. Although it seems plausible that the 3d Ising CFT has *globally* minimal c among all unitary 3d CFTs, this conclusion cannot be reached by studying just a single four-point function of the theory.

Secondly, the minimum over $\mathcal{C}_{\Delta_\sigma}$ in Eq. (2.10) must be computed assuming an additional lower cutoff on Δ_ϵ . The point is that in the region $\Delta_\epsilon < 1$ there exist c minima with $c < c_{3d \text{ Ising}}$. These clearly have nothing to do with the 3d Ising CFT, which is known to have $\Delta_\epsilon \approx 1.412$. In fact, we know of no examples of 3d unitary CFTs with $\Delta_\epsilon < 1$, and we

¹⁰We are aware of the fact that c is not always monotonic under RG-flow, e.g. [13].

suspect that these extra minima are altogether unphysical. In this work we eliminate them by imposing a cutoff $\Delta_\epsilon \geq \Delta_\epsilon^{\text{cutoff}} \approx 1$ (this was already done in producing Fig. 4). The end results are not sensitive to the precise value of this cutoff. In particular, the final spectra do not contain an operator at the cutoff.

From a practical standpoint, c -minimization or p_T -maximization over $\mathcal{C}_{\Delta_\sigma}$ is a linear programming problem which can be attacked numerically on a computer.¹¹ For this purpose, we have implemented a customized version of Dantzig’s simplex algorithm, capable of dealing with the continuous infinity of possible dimensions Δ , and exploiting special structure in the conformal blocks $G_{\Delta,\ell}$. Our algorithm is described in detail in section 6.2. In the next section, we will apply it to p_T -maximization and study the spectrum of the 3d Ising CFT.

3 The 3d Ising Spectrum from c -minimization

In the previous section, we conjectured that the 3d Ising CFT lives at the overall c minimum in Figure 4. In this section, we will focus on the interval of $\Delta_\sigma \in [0.5179, 0.5185]$ containing this minimum. We will perform a high-precision exploration of this interval, including as many as $N = 231$ derivative constraints. This is the highest number of constraints ever used in the analysis of a single bootstrap equation, and is a threefold increase compared to $N = 78$ used in our previous work [1].¹²

We will obtain a number of predictions for the operator dimensions and the OPE coefficients in the 3d Ising CFT. Comparison with previously known results by other techniques is postponed to section 4.

3.1 Δ_σ and c

In Figure 7, we show the c minimization bound computed as a function of Δ_σ in the region near the overall c minimum, for $N = 153, 190, 231$.¹³ The shape of these bounds can be schematically described as follows: on the left an almost-linearly decreasing part, then a “kink” and a short middle part leading to the overall minimum, and finally a monotonically increasing part on the right. When N is increased, the bound is getting visibly stronger on the right, but is essentially stable on the left. As a result the minimum is shifting somewhat upward and to the left, while the kink is barely moving.

In these plots, the middle part of the bounds between the kink and the minimum is shrinking monotonically when N is increased. It looks certain that the minimum is set to merge with the kink in the $N \rightarrow \infty$ limit. According to our conjecture, this merger

¹¹Because Δ can vary continuously, p_T -maximization should more properly be called a *semi-infinite* program, although we will not be careful about this distinction.

¹²In the analysis of CFTs with global symmetry [12], up to 66 constraints were used per each of the three (for $SO(N)$) or six (for $SU(N)$) bootstrap equations present in that case. In [4], up to 55 constraints were used per bootstrap equation. See Appendix A for a discussion of the methods used in these works.

¹³Here and in subsequent plots $N = 231$ data cover a smaller subinterval $\Delta_\sigma \in [0.5180, 0.5183]$.




Figure 7: A lower bound on c in the region of Δ_σ close to the minimum in Figure 4. We plot three bounds, computed with $N = 231$ (dark blue), 190 (lighter blue), 153 (lightest blue). The pink rectangle in the zoomed inlay corresponds to our conservative upper and lower bounds (3.1) for the location of the minimum for $N \rightarrow \infty$, which according to our conjecture should correspond to the 3d Ising CFT.

point should correspond to the 3d Ising CFT. Therefore, the positions of the kink and the minimum at $N = 231$ give upper and lower bounds on Δ_σ and c in the 3d Ising CFT (pink rectangle in the zoomed inlay in Figure 7):

$$\Delta_\sigma = 0.518154(15), \quad c/c_{\text{free}} = 0.946534(11). \quad (3.1)$$

For simplicity, we give symmetric error bars around the center of the middle part of the $N = 231$ bound. However, from the way the middle part of the bounds is shrinking mostly from the right, we can foresee that the true 3d Ising CFT values lie probably closer to the left kink.

3.2 Extracting the Spectrum and OPE Coefficients in $\sigma \times \sigma$

The c -minimization bounds in Figure 7 actually contain much more information than plotted there. For each point saturating those bounds, there is a *unique* unitary four-point function $\langle \sigma\sigma\sigma\sigma \rangle$ which solves the corresponding N crossing constraints. It is very interesting to study how the spectrum of operators appearing in the conformal block decomposition of this four-point function, and their OPE coefficients, depend on Δ_σ . This will be done in the next section for the leading scalar, and in the following sections for the higher states.

We should stress that no additional computation is needed to extract the solution to crossing corresponding to the minimal c . The minimization algorithm starts somewhere inside the allowed region (i.e. above the bound) and performs a series of steps. Each step replaces a solution to crossing by another solution with a strictly smaller c , thus

moving towards the boundary. After many steps (tens of thousands), c stops varying appreciably, and the algorithm terminates. We thus obtain both the minimal c value and the corresponding solution to crossing. Empirically, we observe that the spectrum and the OPE coefficients change a lot from one step to the other in the initial phases of the algorithm, while in the end they stabilize to some limiting values, which depend just on Δ_σ and the value of N we are working at. These limiting values depend smoothly on Δ_σ , except for some interesting rearrangements in the higher states happening near the 3d Ising point, which will be discussed below. The smoothness of these limits is by itself also evidence for their uniqueness. As mentioned above, to reach the boundary, the simplex method must perform tens of thousands of small steps, different for even nearby values of Δ_σ . The absence of uniqueness would be detectable by jittering in the plots below, but we observe no such jittering.

After these general preliminary remarks, let us explore the spectrum and the OPE coefficients corresponding to the c bounds in Figure 7.

3.3 The Leading \mathbb{Z}_2 -Even Scalar ϵ

In this section, we are interested in the leading scalar operator $\epsilon \in \sigma \times \sigma$. In Figure 8 we plot its dimension Δ_ϵ as a function of Δ_σ .




Figure 8: The dimension of the leading scalar $\epsilon \in \sigma \times \sigma$, from the four-point functions realizing the minimal c bounds in Figure 7, with the same line color assignment. The pink rectangle has the same horizontal extension as in Figure 7; its vertical extension gives our prediction (3.2) for Δ_ϵ in the 3d Ising CFT.

The curves in Figure 8 look qualitatively similar to the Δ_ϵ bound in Figure 1, although it should be kept in mind that they have been computed by a different method. In Figure 1, we were maximizing Δ_ϵ , while in Figure 8 we minimized c and extracted Δ_ϵ corresponding

to this minimum. Nevertheless, as discussed in section 2.4, the two methods give very close results near the 3d Ising point, so here we are using the c -minimization method.

The plots in Figure 8 have a narrowly localized kink region, which keeps shrinking as N is increased. Just as we used the c bounds to give predictions for Δ_σ and c , we can now use Figure 8 to extract a prediction for Δ_ϵ . The upper and lower bounds are given by the pink rectangle in the zoomed inlay. Notice that the horizontal extension of this rectangle is exactly the same as in Figure 7—the changes in slope happen at precisely the same Δ_σ 's as for the c bounds. This is not surprising since c and Δ_ϵ enter as variables in the same bootstrap equation, and any non-analyticity in one variable should generically have some reaction on the others. On the other hand, the vertical extension of the rectangle gives our confidence interval for Δ_ϵ :

$$\Delta_\epsilon = 1.41267(13). \quad (3.2)$$

In Figure 9 we repeat the same exercise for the ϵ 's OPE coefficient. Again, the horizontal extension of the pink rectangle is the same as in Figure 7, while its vertical extension gives our prediction for the OPE coefficient:

$$f_{\sigma\sigma\epsilon}^2 = 1.10636(9). \quad (3.3)$$




Figure 9: Same as the previous figure, for the squared OPE coefficient $f_{\sigma\sigma\epsilon}^2$.

In this paper we normalize the OPE coefficients and conformal blocks in the same way as in [1]. For scalars in 3d this normalization coincides with the most commonly used normalization of OPE coefficients through the three-point function:

$$\langle \mathcal{O}_1(x_1) \mathcal{O}_2(x_2) \mathcal{O}_3(x_3) \rangle = \frac{f_{123}}{x_{12}^{\Delta_1+\Delta_2-\Delta_3} x_{13}^{\Delta_1+\Delta_3-\Delta_2} x_{23}^{\Delta_2+\Delta_3-\Delta_1}}, \quad (3.4)$$

where all scalars are assumed to have a unit-normalized two-point function. For example, in Mean Field Theory we have $f_{\phi\phi\mathcal{O}} = \sqrt{2}$ where ϕ and $\mathcal{O} = \phi^2/\sqrt{2}$ are unit-normalized.

3.4 Higher Scalars: General Features

Let us now take a look at the higher scalars in the four-point function minimizing c . In Figure 10, we show how the scalar dimensions and OPE coefficients vary in the interval $\Delta_\sigma \in [0.516, 0.52]$. These plots correspond to the c bound at $N = 105$ from Figure 5. In Figure 11, we zoom in on the interval close to the 3d Ising point. Here we are using higher N data corresponding to the c bounds in Figure 7.




Figure 10: Scalar operators with dimensions with $\Delta \leq 13$ and their squared OPE coefficients (matching color and numbering). Here we are using the same $N = 105$ data as in Figure 5 (right). The squared OPEs are plotted multiplied by 4^{Δ_i} to somewhat correct for the disparity in size; this is also the natural normalization from the point of view of the OPE convergence estimates [9].




Figure 11: Scalar operators with dimensions with $\Delta \leq 13$ and their squared OPE coefficients, from the four-point functions realizing the c bounds in Figure 7: $N = 231$ (darker shades), 190 (lighter shades), 153 (lightest shades).

In these plots we are showing all scalars with $\Delta \leq 13$. The blue (no.1) curves correspond to the leading scalar ϵ . As we have seen in the previous section, its dimension and OPE coefficient vary slowly; at the scale of this plot they look essentially constant. It is then

shocking to see how much the higher scalar curves are varying on the same scale. The most salient properties of these curves can be summarized as follows:

1. The higher scalar operator dimensions and OPE coefficients vary mildly, except at the 3d Ising point, where they experience rapid, near-discontinuous changes. The transition region where these changes happen is shrinking as N is increased (this is especially noticeable in the OPE coefficient plot in Figure 11).
2. The effect of the above changes is just to shift the higher scalar spectrum by one operator dimension up, *à la* Hilbert’s infinite hotel. The operator marked no.2 (red) of dimension ≈ 3 below the 3d Ising point disappears. As a result the higher scalar spectrum and the OPE coefficients above the 3d Ising point (and past the transition region) are the same as below, minus the disappearing operator.
3. The OPE coefficient of the disappearing operator no.2 tends to zero approaching the 3d Ising point. This property may not be obvious from the shown plots, as it’s obscured by the presence of the transition region, but we believe that it should become exact in the limit $N \rightarrow \infty$. What we mean is that the red (no.2) OPE coefficient curve in Figure 10 is heading towards zero before entering the transition region, at which point it’s shooting up to reconnect with the green (no.3) curve. From Figure 11 we can see that the minimum value this curve reaches before shooting up becomes smaller and smaller as N becomes larger. In Figure 12 we give an idealized sketch showing what the above plots should look like at a much larger N than considered here.




Figure 12: An idealized rendering of the likely scalar operator behavior at a much larger N . Between the dashed lines is the transition region. We believe that all operator dimensions should go to some slowly varying limits outside of this ever-shrinking region. In particular the red (no.2) operator which is currently showing larger variation than say the green (no.3) operator should also become slowly varying. All OPE coefficients should also become slowly varying except the red (no.2) OPE coefficient going to zero when approaching the 3d Ising point from below. Apart from the red (no.2) operator disappearing, there is one-to-one correspondence between the spectrum and OPE coefficients above and below the transition region. In the $N \rightarrow \infty$ limit, the spectrum and OPE coefficients should vary continuously (although not necessarily smoothly) across the 3d Ising point, with one operator disappearing.

The first hints of property 1 were noticed in our previous work [1],¹⁴ where we presented an *upper bound* on the dimension of the subleading scalar $\Delta_{\ell'}$, fixing Δ_{ℓ} to its maximal allowed value, i.e. the boundary in Figure 1. That upper bound¹⁵ showed strong variation at the 3d Ising point, just as curve no.2 does. As we now realize, that upper bound was in fact more than simply a bound—it was showing the unique value of $\Delta_{\ell'}$ allowed at the boundary.

Although Properties 2 and 3 are exhibited here for the first time, in retrospect they are in fact very natural and connected to the very existence of the 3d Ising kink. Imagine approaching the 3d Ising point along the boundary of the c lower bound in Figure 7. All along the bound we have a solution to crossing. If everything in this solution varies smoothly, we can analytically continue the solution beyond the 3d Ising point. Yet we know that this is impossible, since the bound shows a kink there. So something must happen which invalidates the analytically continued solution. The simplest obstruction is if some $p_{\Delta,\ell}$ hits zero at the 3d Ising point, so that the analytic continuation has negative $p_{\Delta,\ell}$ and violates unitarity. Property 3 means that such an obstruction is indeed encountered when approaching 3d Ising from below, as one $p_{\Delta,\ell=0}$ hits zero. We will see in section 3.6 that an obstruction of the same type occurs when approaching the 3d Ising CFT from above, except that in this case the operator whose OPE coefficient hits zero has $\ell = 2$.

Property 2 implies a practical way to extract the 3d Ising CFT spectrum—it is given by the operators dimensions which are present on both sides of the transition region. The operator no.2 (red curve) on the left of the transition region is thus excluded, since it decouples at the 3d Ising point. Looking at Figure 11 and applying this rule, we expect that the second \mathbb{Z}_2 -even scalar after ϵ has dimension ≈ 4 , since this dimension is present both below (green no.3) and above (red no.2) the transition region. In the next section we will be able to determine its dimension and the OPE coefficient rather precisely. A third \mathbb{Z}_2 -even scalar ϵ'' appears at dimension ≈ 7 , and a fourth ϵ''' at ≈ 10.5 . These estimates (especially the one for ϵ''') are preliminary, since the curves corresponding to these operators show non-negligible variation for values of N in Figure 11. In particular, we prefer not to assign error bars to them.

Although we do not show scalars of $\Delta > 13$ in Figure 11, we checked that the same qualitative behavior continues. In particular, the operator dimensions continue shifting by one operator across the transition region, and no additional operator decoupling is observed, so that the red no.2 line below the 3d Ising point is the only one which decouples. The higher in dimension one goes, the stronger is variation with N . This loss of sensitivity is expected in view of the exponential decoupling of high dimension operators in a four-point function of low-dimension operators [9]. A natural way to boost sensitivity to high exchanged dimensions might be to raise the dimension of the external operators, by considering e.g. the four-point function $\langle \epsilon \epsilon \epsilon \epsilon \rangle$. It would be important to try this in the future.

¹⁴And even earlier in the two-dimensional case in [14].

¹⁵As well as an analogous bound on the dimension of the subleading spin 2 operator.

3.5 The Second \mathbb{Z}_2 -Even Scalar ϵ'

Having discussed the general properties of the 3d Ising scalar spectrum in the previous section, here we will focus on the second \mathbb{Z}_2 -even scalar ϵ' , which is the operator of dimension ≈ 4 present both below (green no.3) and above (red no.2) the transition region in Figure 11. In Figure 13 we give a vertical zoom of the part of Figure 11 where the green no.3 and the red no.2 lines approach each other.




Figure 13: The zoom on the part of the scalar spectrum relevant for the extraction of the ϵ' parameters. The operator colors and numbering are the same as in Figure 11. Note that compared to Figure 11, the OPE coefficients are here plotted without the 4^Δ factor.

We fixed the horizontal extension of the pink rectangles in these plots equal to the 3d Ising Δ_σ range previously determined in sections 3.1 and 3.3. We see that this range falls on the part of the red no.2 plateau which is quite converged. The vertical extension of these rectangles then gives our best estimates of the ϵ' dimension and the OPE coefficient:

$$\Delta_{\epsilon'} = 3.8303(18), \quad f^2_{\sigma\sigma\epsilon'} = 0.002810(6). \quad (3.5)$$

In contrast to the red no.2 plateau, the green no.3 curves to the left of the 3d Ising point are still changing significantly with N . As explained in the previous section, we expect that in the $N \rightarrow \infty$ limit the green no.2 curves will reach a limiting plateau continuously connecting to the red no.2 plateau. Although this has not happened yet on the scale of Figure 13, the tendency is clearly there.

3.6 Spin 2 Operators

In this section we analogously consider the $\ell = 2$ operators in the $\sigma \times \sigma$ OPE. In Figure 14 we show the $\ell = 2$ spectrum and OPE coefficients at $N = 105$ in a wider range, and in Figure 15 we zoom in at the transition region at higher N .

We see from these plots that there are many similarities in the general features of $\ell = 0$ and $\ell = 2$ spectra. The lowest operator in the spectrum is now the stress tensor, its




Figure 14: Spin 2 operators of dimension $\Delta < 16$ and their OPE coefficients, at $N = 105$.

dimension is 3 and independent of Δ_σ , Its OPE coefficient varies slightly, according to the c -minimization bounds shown above, but on this scale this variation is not noticeable.

For the higher spectrum, the analogues of Properties 1,2,3 from section 3.4 are true, with one difference: the red no.2 operator now decouples approaching the 3d Ising point from the right instead of from the left. The fact that its OPE coefficient tends to zero in this limit (and for $N \rightarrow \infty$) is even more evident here than it was for its cousin in section 3.4. As promised, the existence of this decoupling operator provides an obstruction for the analytic continuation across the kink, approaching it from above.




Figure 15: Spin 2 operators of dimension $\Delta < 16$ and their OPE coefficients, at $N = 231, 190, 153$.

Leaving out the decoupling operator and interpolating the plateaux above and below the transition region in Figure 15, we get an approximate $\ell = 2$ spectrum in the 3d Ising CFT. Apart from the stress tensor, it should contain operators of dimensions $\approx 5.5, 8, 10, 13, \dots$

We will now determine the parameters of the first of these subleading operators, call it T' , similarly to how we fixed ϵ' in section 3.5. In Figure 16 we zoom in on the region of near level-crossing between the red no.2 and the green no.3 curves. The horizontal extent of the

pink rectangle coincides with the 3d Ising Δ_σ confidence interval. We are again lucky in that this interval falls on the part of the red plateau which looks reasonably converged. So we determine the vertical extension of the rectangles by fitting in the $N = 231$ curves, and obtain the following estimates:

$$\Delta_{T'} = 5.500(15), \quad f_{\sigma\sigma T'}^2 = 2.97(2) \times 10^{-4}. \quad (3.6)$$




Figure 16: The zoom on the region relevant for the extraction of the T' operator parameters. The operator colors and numbering are the same as in Figure 15. Note that the OPE coefficients are plotted without the 4^Δ factor.

From Figures 13 and 16 it can be seen clearly that the left (right) end of the Δ_σ confidence interval coincides with the location of the operator jumps in the $\ell = 0$ ($\ell = 2$) sector. Thus looking at these jumps gives an equivalent method to localize the 3d Ising point.

3.7 Higher spins

The $\sigma \times \sigma$ OPE is also expected to contain infinitely many primary operators of spin $\ell = 4, 6, 8, \dots$. It would be of course interesting to learn something about their dimensions from the bootstrap. Particularly interesting are the operators of the smallest dimension for each spin, let's call them C_ℓ . Being an interacting CFT, the critical Ising model cannot contain conserved higher spin currents [15]. Thus operators C_ℓ should have a positive anomalous dimension:

$$\Delta(C_\ell) = \ell + d - 2 + \gamma_\ell, \quad \gamma_\ell > 0. \quad (3.7)$$

The dimension of C_4 is known with some precision [16]:

$$\Delta(C_4) = 5.0208(12). \quad (3.8)$$

Dimensions of higher C_ℓ are not accurately known, although at the Wilson-Fisher fixed point in $d = 4 - \epsilon$ [17]

$$\gamma_\ell = \frac{\epsilon^2}{54} \left(1 - \frac{6}{\ell(\ell+1)} \right) + O(\epsilon^3). \quad (3.9)$$

There are also two general results about the sequence of γ_ℓ , which are supposed to be valid in any CFT. Firstly, the large ℓ limit of γ_ℓ has to be equal twice the anomalous dimension of σ [18, 19]:¹⁶

$$\lim_{\ell \rightarrow \infty} \gamma_\ell = 2\gamma_\sigma \equiv 2\Delta_\sigma - 1. \quad (3.10)$$

The asymptotic rate of approach to the limit is also known, see [18, 19].

Secondly, we have “Nachtmann’s theorem” [21], which says that the sequence γ_ℓ is monotonically increasing and upward convex. This remarkable result is on a somewhat less solid footing than (3.10). As was emphasized in its recent discussion in [19], its derivation uses the polynomial boundedness of a certain scattering amplitude, which is still conjectural (although plausible) at present. Depending on the degree of the polynomial boundedness, Nachtmann’s theorem may hold not for all ℓ but in the range $\ell \geq \ell_0$.

It would be therefore very interesting to determine γ_ℓ in the 3d Ising CFT using the bootstrap, and see if they satisfy the above two general properties.

In principle, the determination of γ_ℓ for $\ell \geq 4$ using the c -minimization is as straightforward as the determinations of $\ell = 0, 2$ operator dimensions discussed in the previous sections. The extremal c -minimization spectra that we obtain do contain higher spin operators, and we can easily identify the lowest operator for each spin.

In practice, however, this procedure for $\ell \geq 4$ turns out to be numerically somewhat less stable than for $\ell = 0, 2$. This must be somehow related to the fact that the anomalous dimensions γ_ℓ are all expected to be small yet nonzero. Since conformal blocks of spin $\ell \neq 0$ operators are continuous in the limit of Δ approaching the unitarity bound, it’s not easy for our algorithm to distinguish an exactly conserved operator from one with a small anomalous dimension. Moreover, in this case, $d = 2$ does not provide any guidance as there the higher spin operators *are* conserved and our algorithm has no difficulty in reconstructing the $\ell > 2$ spectrum (see section 5.1 below).

In spite of these numerical problems, our calculations do show that the higher spin currents acquire positive anomalous dimensions. We managed to extract anomalous dimensions up to spin $\ell \simeq 40$. Although precision needs to be improved, the extracted values are largely consistent with the Callan-Gross relation (3.10). Nachtmann’s theorem also seems to hold, in the full range of $\ell \geq 4$. Our value of $\Delta(C_4)$ is roughly consistent with (3.8).

These preliminary results are encouraging, and we give them here since they may stimulate further work on Nachtmann’s theorem. Nevertheless, we prefer to postpone detailed plots and precision determinations until we have the higher spin sector under better control.

¹⁶This is sometimes called the Callan-Gross relation since it was first noticed in perturbation theory in [20].

4 Comparison to Results by Other Techniques

In the previous section, we used the conformal bootstrap method together with our c -minimization conjecture to determine several operator dimensions and OPE coefficients of the 3d Ising CFT. We will now compare our results to prior studies of the 3d Ising model at criticality, by other techniques.

Operator dimensions and critical exponents

One well studied class of universal quantities characterizing the critical point are the critical exponents.¹⁷ They have simple expressions in terms of the CFT operator dimensions. In particular, the well-known critical exponents η , ν , ω can be expressed via the dimensions of σ , ϵ , and ϵ' from the formulas:

$$\Delta_\sigma = 1/2 + \eta/2, \quad \Delta_\epsilon = 3 - 1/\nu, \quad \Delta_{\epsilon'} = 3 + \omega. \quad (4.1)$$

In Table 2 we quote some notable prior determinations of these exponents. The first two determinations from [22] are by field theoretic (FT) techniques (the ϵ -expansion and fixed-dimension expansion). For η and ν , these are an order of magnitude less precise than the best available determination from the lattice: the Monte Carlo simulations (MC) and the high-temperature expansion (HT). It is again the MC which provides the best estimate of ω , followed by FT, and HT.

ref	year	Method	ν	η	ω
[22]	1998	ϵ -exp	0.63050(250)	0.03650(500)	0.814(18)
[22]	1998	3D exp	0.63040(130)	0.03350(250)	0.799(11)
[23]	2002	HT	0.63012(16)	0.03639(15)	0.825(50)
[24]	2003	MC	0.63020(12)	0.03680(20)	0.821(5)
[25]	2010	MC	0.63002(10)	0.03627(10)	0.832(6)
this work			0.62999(5)	0.03631(3)	0.8303(18)

Table 2: The most precise prior determinations of the leading critical exponents, compared to our results.

In the same table we give the values of η , ν , ω obtained via (4.1) using the values of operator dimensions from the previous section. Comparing our values with the recent precise MC determinations by Hasenbusch [25], we see that the agreement is very good, our results being factor 2 - 3 more precise. The agreement with the older MC results [24] is not as good, with about 2σ tension in every exponent. Our results clearly favor [25] over [24].

Subleading correction-to-scaling exponents

Another method for computing the critical exponents is the functional renormalization group (FRG), in its various incarnations. For the leading exponents η , ν , ω , it is not

¹⁷Another such class are the amplitude ratios. These are related to IR-dominated properties of RG flows produced when the theory is perturbed away from the critical point. Unlike the critical exponents, the amplitude ratios cannot be easily computed in terms of the CFT parameters.

ref	year	FRG variant	ω_2	ω_3	ν	η	ω
[28]	1984	scaling field (13 ops)	1.67(11)	—	0.626(9)	0.040(7)	0.855(70)
[29]	1997	∂^0 , sharp cutoff	2.8384	5.1842	0.6895	0	0.5952
[30]	2002	∂^0 , optimized cutoff	3.180	5.912	0.6496	0	0.6557
[31]	2003	∂^0 , quartic cutoff	3.048	5.63	0.6604	0	0.6285
		∂^0 , background field	3.6845	7.038	0.6260	0	0.7622
this work			≈ 4	≈ 7.5	0.62999(5)	0.03631(3)	0.8303(18)

Table 3: Functional renormalization group determinations of the subleading correction-to-scaling exponents, compared to our results. For completeness and ease of comparison, we cite also the results for the leading exponents obtained in the same studies.

currently competitive with the methods listed in Table 2 (see [26, 27] for state-of-the-art studies). On the other hand, an advantage of this method is that it can compute the subleading correction-to-scaling exponents. Here we will be particularly interested in the subleading exponents $\omega_{2,3}$ related to the dimensions of the higher \mathbb{Z}_2 -even scalar operators ϵ'' and ϵ''' :

$$\Delta_{\epsilon''} = 3 + \omega_2, \quad \Delta_{\epsilon'''} = 3 + \omega_3. \quad (4.2)$$

In Table 3 we collect several available FRG determinations of $\omega_{2,3}$ and compare them with our results. The oldest calculation is by the “scaling field” method [28], which performs an RG flow in a local operator basis truncated to finitely many operators and optimized to minimize the truncation effects. The largest bases had relatively few (≤ 13) operators, but unlike the calculations discussed below they included some operators with two and four derivatives. This method gives the leading critical exponents in agreement, within cited errors, with the more precise determinations from Table 2. On the other hand, their value of ω_2 is much smaller than our estimate $\omega_2 \approx 4$.

Then there is a group of more recent calculations which truncate the flowing Lagrangian to the standard kinetic term plus terms with an arbitrary number of ϕ ’s but no derivatives; these are marked as ∂^0 in the table. By the nature of such a truncation, all these calculations give $\eta = 0$. The other exponents vary significantly depending on the type of cutoff used to regularize the RG flow. The background field cutoff gives ν and ω closest to our determinations; it also gives the values of $\omega_{2,3}$ closest to our estimates.

In view of the remaining discrepancies, it would be interesting to upgrade the recent FRG calculations of the subleading exponents by increasing the derivative truncation order to ∂^2 and ∂^4 . While such calculations have been performed for the leading exponents and noticed to improve their estimates [26, 27], we are not aware of any results for $\omega_{2,3}$.

OPE coefficient $f_{\sigma\sigma\epsilon}$

Hardly anything is known about the OPE coefficients in the 3d Ising CFT. Some universal couplings were studied by Henkel [32], but as we will now explain they are not equal to the flat space OPE coefficients. He studied a 2d Hamiltonian with a quantum critical point known to belong to the 3d Ising universality class. The Hamiltonian was diagonalized on a family of square $N \times N$ lattices with periodic boundary conditions ($N \leq 5$). Continuum

limit quantities were extracted by applying finite-size scaling analysis to matrix elements of lattice versions of ϵ and σ operators sandwiched between the vacuum and the lowest nontrivial \mathbb{Z}_2 -even and \mathbb{Z}_2 -odd eigenstates:¹⁸

$$\langle 0|\epsilon|+\rangle \quad \text{and} \quad \langle 0|\sigma|-\rangle. \quad (4.3)$$

This gave reasonable, although not very precise, estimates for the continuum operator dimensions: $\Delta_\sigma = 0.515(9)$ and $\Delta_\epsilon = 1.42(2)$. He then considered universal quantities related to the properly normalized matrix elements:

$$C_{\sigma\sigma\epsilon} \propto \langle -|\sigma|+\rangle, \quad C_{\sigma\epsilon\sigma} \propto \langle -|\epsilon|-\rangle. \quad (4.4)$$

Are these in any way related to $f_{\sigma\sigma\epsilon}$? In 2d, the answer would be yes, but in 3d it's no, for the following reason. Using periodic square lattices means that the critical 3d theory is put on $T^2 \times \mathbb{R}$. Since this geometry is not Weyl-equivalent to flat 3d space, there is no way to relate the C 's to the flat space OPE coefficients. A priori, the two C 's do not even have to be equal; and in fact [32] finds unequal values:

$$C_{\sigma\sigma\epsilon} \approx 0.85, \quad C_{\sigma\epsilon\sigma} \approx 1.10. \quad (4.5)$$

In order to measure $f_{\sigma\sigma\epsilon}$, one would have to put the theory on $S^2 \times \mathbb{R}$. Via the radial quantization state-operator correspondence, the states on S^2 are obtained by the flat-space operators inserted at the origin and at infinity. The lowest states $|-\rangle$ and $|+\rangle$ can then be identified with $|\sigma\rangle$ and $|\epsilon\rangle$, and we would have

$$C_{\sigma\sigma\epsilon} = C_{\sigma\epsilon\sigma} = f_{\sigma\sigma\epsilon} \quad (\text{on } S^2 \times \mathbb{R}). \quad (4.6)$$

In connection with this we would like to point out a recent paper [33] which measures Δ_σ via a MC simulation of the statistical mechanics 3d Ising model in the $S^2 \times \mathbb{R}$ geometry, approximating S^2 by a discretized icosahedron.¹⁹ It would be interesting to measure $f_{\sigma\sigma\epsilon}$ by this or another technique, so that we would have something to compare to.

Central charge c

The prior knowledge of c can be summarized by the ϵ -expansion formula [35–38]:

$$c/c_{\text{free}} = 1 - \frac{5}{324}\epsilon^2 + \dots \quad (4.7)$$

Preliminary bootstrap determinations of c in 3d were given in [1, 3]. Moreover, Ref. [3] checked the above formula via the conformal bootstrap analysis in a fractional number of spacetime dimensions $d = 4 - \epsilon$. Good agreement was observed for $\epsilon \lesssim 0.3$, while in 3d ($\epsilon = 1$) corrections due to the unknown $O(\epsilon^3)$ and higher terms must be significant.

It would be interesting if the 10^{-5} precision of our c determination in section 3.1 could be approached or matched by any other technique.

¹⁸Henkel refers to $|-\rangle$ and $|+\rangle$ as $|\sigma\rangle$ and $|\epsilon\rangle$, but as will see below this notation may lead to a confusion in his geometry.

¹⁹An earlier reference [34] measured Δ_σ and Δ_ϵ approximating S^2 by a discretized cube.

5 2d Checks

So far in this paper we were applying the bootstrap methods and the c -minimization conjecture to extract information about the 3d Ising CFT. Since it is known that the 3d Ising CFT is a member of the family of the Wilson-Fisher CFTs interpolating between 2 and 4 dimensions, it would be interesting to apply the same techniques in $2 \leq d < 4$ to learn more about this family, which so far has been studied mostly using the ϵ -expansion [39].²⁰ First steps in this direction were made in our paper [3], where we used the position of the Δ_ϵ -maximization kink as a function of $2 \leq d < 4$ to extract Δ_σ , Δ_ϵ , and (from the extremal solution) c .

Postponing a more detailed study of fractional dimensions to the future, in this section we will discuss in some detail the other integer case, $d = 2$. This case has already played a very stimulating role in the development of our techniques. It was first observed in [41] that the 2d Δ_ϵ -maximization plot has a kink whose position could be identified with the exactly known operator dimensions of the 2d Ising CFT $\Delta_\sigma = 1/8$, $\Delta_\epsilon = 1$. The sharp variation of the subleading scalar dimension near the kink was also first noticed in 2d [14]. Both these facts turned out to have close analogues in 3d.

In this section, we will go beyond the previous 2d studies and apply our method to extract the low-lying $\sigma \times \sigma$ spectrum and OPE coefficients in the neighbourhood of the 2d Ising CFT. At the 2d Ising point we will be able to improve on the accuracy of previous results [2] and determine Δ_σ itself from the position of the kink. In a region to the right of the kink our method closely reconstructs a certain family of solutions to crossing symmetry related to the 2d minimal models. In particular, the spectrum of quasiprimaries will respect the Virasoro symmetry, even though our bootstrap equations have only global $SL(2, \mathbb{C})$ symmetry built in from the start.

Apart from checking the c -minimization method against the exact solution, the 2d study will also shed some light on the operator decoupling phenomenon observed in 3d when approaching the Ising point. As we will see, an analogous decoupling happens in 2d. While in 3d the reasons behind the decoupling remain unknown, in 2d it can be given a natural explanation, in terms of null states present in the 2d Ising CFT.

5.1 Spectrum Extraction at Fixed $\Delta_\sigma = 1/8$

In section 3, spectrum determination in the 3d Ising CFT consisted of two steps. First, a range for $\Delta_\sigma^{\text{3d Ising}}$ was identified near the minimum of the $c(\Delta_\sigma)$ curve. Second, spectra corresponding to the solutions in this range were looked at and interpreted under the assumptions of convergence, operator rearrangement and decoupling.

In 2d, Δ_σ is known exactly, so it is possible to eliminate the first step and study the second step in isolation [2]. We fix $\Delta_\sigma = 1/8$ and maximize Δ_ϵ , which gives a value very

²⁰See also [40] for a study based on the functional renormalization group.

close to the exact value $\Delta_\epsilon = 1$.²¹ We then extract the low-lying spectrum and the OPE coefficients corresponding to the maximal Δ_ϵ .




Figure 17: The extremal 2d spectrum extracted at $\Delta_\sigma = 1/8$. This plot was produced using our code at $N = 120$. Such results were first obtained in [2] at $N = 60$, using the dual method.

The results (thin red bars) are shown and compared to the exact 2d Ising data (wide blue bars) in Figure 17, for spins $\ell \leq 18$ and dimensions $\Delta \leq 20$. Let us first discuss the exact data. The blue bars are centered at the positions of the \mathbb{Z}_2 -even quasiprimaries $\mathcal{O}_i \in \sigma \times \sigma$ in the 2d Ising CFT, the height of the bars given by the squared OPE coefficients.²² All these

²¹ $\Delta_\epsilon^{\max} = 1.000003$ for $N = 60$ [2].

²²The OPE coefficients were obtained by expanding the exactly known four-point function $\langle \sigma\sigma\sigma\sigma \rangle$ into $SL(2, \mathbb{C})$ conformal blocks. When there are several quasiprimaries with the same dimension and spin, they

operators are obtained from $\mathbb{1}$ and ϵ by acting with higher Virasoro generators L_{-n} , $n \geq 2$, and keeping track of the null state conditions specific to the 2d Ising CFT. For example, the leading twist trajectory $\Delta = \ell$ (the highest diagonal) are the Virasoro descendants of the identity. Four and eight units above it lie the operators obtained by acting with $L_{-2}\bar{L}_{-2}$ and $L_{-4}\bar{L}_{-4}$. Twist one operators are the ϵ and its Virasoro descendants; this “ ϵ -trajectory” has a gap for $\ell = 2$ because the $\epsilon = \phi_{1,2}$ is degenerate on level 2, and the state in question is null. For the same reason, the ϵ -trajectory has daughters at 8, 12, 16... but not 4 units higher.

Now consider the red bars, whose position and height show the extremal spectrum and the OPE coefficients extracted by our algorithm at $N = 120$. The agreement for the first three operator trajectories is almost too good. The fourth trajectory can also be divined, although not as cleanly, at the positions where the numerical spectrum has several nearby operators. For even higher twists the sensitivity of our algorithm becomes insufficient, and the numerical spectrum no longer correlates with the exact one. (When N is increased, the sensitivity loss threshold is pushed to larger and larger twists.)

One important lesson from this plot is that for an operator to be numerically extractable, it better have a sizable OPE coefficient: the four extracted trajectories consist of the operators with the largest coefficients. To reach this conclusion it’s important to normalize OPE coefficients using the natural normalization of [9].

5.2 c -Minimization

In the previous section we set Δ_σ to the precisely known 2d Ising CFT value $\frac{1}{8}$. In this section we will instead mimic more closely our 3d study, performing a c -minimization scan in the neighborhood of the 2d Ising point.

In Figure 18 we show the lower bound for c in 2d, using 153 components. Just as in the 3d case (see the discussion in section 2.4), we have to impose an additional constraint $\Delta_\epsilon \geq \Delta_\epsilon^{\text{cutoff}}$ to eliminate spurious minima far away from the Ising point. The end results are insensitive to $\Delta_\epsilon^{\text{cutoff}}$ as long as it is sufficiently large. In this study we found $\Delta_\epsilon^{\text{cutoff}} \approx 0.5$ suffices.²³

The inset in Figure 18 shows that by 153 components the minimum seems to have converged, to quite high precision, to the expected values $(\Delta_\sigma, c) = (\frac{1}{8}, \frac{1}{2})$. To be precise, the c minimum in this plot is 0.499999. Thus in 2d we can determine c with 10^{-6} precision at $N = 153$, while in 3d we only had 10^{-5} precision at $N = 231$ —the bootstrap analysis seems to converge faster in two dimensions.

A precious feature of the 2d situation is that, unlike in the 3d case, we have a very good guess not only for the physical meaning of the point realizing the overall minimum of c —the

are lumped together by summing their OPE coefficients squared.

²³The first study of the c lower bound in two dimensions was done in [42], section 6.3.2. A peculiarity of the 2d case is that for $\Delta_\epsilon^{\text{cutoff}}$ close to the unitarity bound, the lower bound on c disappears altogether, allowing solutions to crossing with an arbitrarily small c .




Figure 18: *Blue curve:* a lower bound on c in two dimensions similar to Fig. 7 but with only $N = 153$ components. In the inlay we show the region near $\Delta_\sigma = \frac{1}{8}$, from which it's clear c is minimized, to high precision, at the exact 2d Ising value $c = \frac{1}{2}$. *Black crosses:* exact interpolating solution discussed in section 5.3 below. The interpolating solution is defined for any $\frac{1}{8} \leq \Delta_\sigma \leq \frac{1}{2}$. The positions of the crosses are chosen arbitrarily to guide the eye; they do not coincide with the positions of the unitary minimal models, of which only two are present in the shown range: $\mathcal{M}(4, 3)$ at $\Delta_\sigma = \frac{1}{8}$ and $\mathcal{M}(5, 4)$ at $\Delta_\sigma = \frac{1}{5}$.

2d Ising CFT—but also for the part of the c bound at $\Delta_\sigma > \Delta_\sigma^{2d \text{ Ising}}$. This part of the curve very likely corresponds to the analytic solution to crossing symmetry interpolating between the unitary minimal models, constructed in [43] and reviewed in section 5.3 below. The black crosses in Figure 18 show that the central charge of the interpolating family follows rather closely our numerical bound. The deviation from the interpolating solution grows somewhat with Δ_σ . At $\Delta_\sigma = \frac{1}{5}$, the interpolating family reaches the second unitary minimal model $\mathcal{M}(5, 4)$ of central charge $\frac{7}{10}$. At this point, the deviation from our $N = 153$ bound is 1.5×10^{-3} . Although this is less impressive than the 10^{-6} agreement at the Ising point, the agreement keeps improving with N , so that the exact interpolating family can plausibly be recovered in the $N \rightarrow \infty$ limit.

As in three dimensions, uniqueness of the extremal solution determines the OPE coefficients and operator dimensions corresponding to the minimal value of c as a function of Δ_σ . The scalar and spin 2 spectra, as well as the OPE coefficients, are plotted in Figures 19 and 20, respectively. We observe the familiar dramatic re-arrangements of the operators at $\Delta_\sigma \sim 1/8$.

In the same figures, we show the dependence of the low-lying operator dimensions and OPE coefficients of the interpolating family for $\Delta_\sigma > \frac{1}{8}$. We see that our numerical solution reproduces well the general features of the interpolating family. In Figure 19, the first, second, and fourth scalars agree very well, together with their OPE coefficients. The third scalar is reproduced less precisely, and in the region close to the Ising point it disappears from our numerical spectrum; this must be due to the fact that it has a very small OPE coefficient and so sensitivity to its dimension is reduced. Indeed, as we discuss in section 5.3, in the exact interpolating solution this scalar becomes null and decouples at the Ising point. The tendency of this state to decouple is well reproduced by our numerics.




Figure 19: *Solid lines:* the spectrum of spin 0 operators with $\Delta < 8.5$ and their OPE coefficients, corresponding to the c bound at $N = 153$. Note the anomalously small OPE coefficient of the green operator ($\Delta \sim 5 - 6$) which tends to decouple slightly to the right of the Ising point. *Circles at $\Delta_\sigma = \frac{1}{8}$:* operators and OPE coefficients in the 2d Ising CFT. *Markers at $\Delta_\sigma > \frac{1}{8}$ ($\times, \Delta, \square, \diamond$):* operators and OPE coefficients in the interpolating solution.

Turning to the spin 2 sector in Figure 20, we see that the first four states in the numerical spectrum all correspond to states in the interpolating family.²⁴ The lowest one is the stress tensor. The second state is decoupling at the Ising point, but its OPE coefficient is not as small as the coefficient of the decoupling scalar, so it is captured well by the numerics.




Figure 20: *Solid lines:* the spectrum of spin 2 operators with $\Delta < 8.5$ and their OPE coefficients, corresponding to the c bound at $N = 153$. *Circles:* operators and OPE coefficients in the 2d Ising CFT. *Markers:* operators and OPE coefficients in the interpolating solution.

A minor blemish is that one more decoupling spin 2 state (marked with question marks) is altogether missed by our numerics. This is perhaps due to it being close to another state with a larger OPE coefficient. Perhaps the observed slight deviation of the 3rd and the 4th

²⁴A state of dimension ~ 6.5 present only in a small interval around $\Delta_\sigma = 0.128$ and with a tiny OPE coefficient is clearly a numerical artifact.

numerical state dimensions can be explained as a distortion induced by trying to compensate for the absence of the “question mark” state. This story should serve as a warning that it will be difficult to resolve closely spaced high-lying states in the future bootstrap studies. On the other hand including additional correlators, where the OPE coefficient in question is larger, should allow us to compensate for this.

Comparing Figures 19 and 20 with the analogous 3d spectrum plots in sections 3.4 and 3.6, we see several similarities as well as some differences. The two most important similarities are, firstly, rapid spectrum rearrangements at the Ising point and, secondly, the subleading spin 2 operator decoupling when approaching the Ising point from the right. In 3d this decoupling is a mystery, but in 2d it has a natural explanation in terms of the interpolating solution. Maybe also in 3d there exists a continuous family of CFTs saturating the c -minimization bound for $\Delta_\sigma > \Delta_\sigma^{\text{3d Ising?}}$

One difference between 2d and 3d is that in 2d we saw a scalar decoupling when approaching the Ising point from the right, while no such scalar was observed in 3d. But the most noticeable difference is in the behavior of the scalar spectrum to the left of the Ising point. In 2d the lowest dimension scalar bifurcates to the left of the Ising model while in 3d it is the subleading scalar which exhibits this behavior. To clarify this puzzle, we carried out a preliminary spectrum study for c -minimization in fractional $2 < d < 3$. We saw that for such d the spectrum to the left of the Ising point is qualitatively similar to the 3d case. In particular, the subleading scalar on the left does not continuously connect to the leading scalar on the right, but curves up and connects to the subleading scalar on the right, as it does for $d = 3$. This “change of topology” happens even though the extremal spectrum as a whole seems to vary continuously with d .

Our last comment concerns the meaning of the extremal numerical solution to the left of the 2d Ising point. Can one find a family of (perhaps non-unitary) CFTs corresponding to this solution, similar to what happens for $\Delta_\sigma > \frac{1}{8}$? We believe that the answer is no, for the following reason. On the one hand, as can be seen Figure 20, the extremal solution on the left contains a local stress tensor (spin 2, dimension 2 operator). A 2d conformal field theory with a local stress tensor will have a Virasoro symmetry, so let’s see if the rest of the spectrum respects it. Unfortunately, the spectrum to the left of the Ising point is not consistent with Virasoro (unlike to the right where it’s largely consistent apart from the missing “question mark” operator). For example, take the two lowest scalars on the left. In the presence of Virasoro symmetry, we would expect to see their spin 2 descendant on level 2, and yet there are no states of such dimension in Figure 20. Perhaps these states are null? However it’s easy to check that the dimension of these scalars is inconsistent with them being degenerate on level 2, given the central charge from Figure 18. Thus we really have a contradiction, and any CFT interpretation of the extremal solution on the left is excluded. This solution must thus be unphysical.

5.3 Interpolating Between Minimal Models

Ref. [43] constructed a family of analytic solutions to crossing which interpolate between the unitary minimal models $\mathcal{M}(m+1, m)$. More precisely, they found a family of crossing symmetric four-point functions $g^{(\Delta_\sigma)}(z, \bar{z})$ with positive conformal block coefficients for all $\frac{1}{8} \leq \Delta_\sigma \leq 1$. These correlators are not realized in a fully unitary CFT unless Δ_σ sits precisely at a minimal model value (this should be visible as a breakdown of positivity in other correlators). However, these correlators do satisfy all the constraints imposed in this work.

As discussed in the previous section, the family of [43] likely saturates the c -minimization for all $\frac{1}{8} \leq \Delta_\sigma \leq \frac{1}{2}$. It also likely saturates the Δ_ϵ -maximization bound in the same range.²⁵

Because the correlators $g^{(\Delta_\sigma)}(z, \bar{z})$ provide an important reference for our results, let us review their construction [43]. Minimal model primaries $\phi_{r,s}$ are labeled by integers r, s . Let us take $\sigma = \phi_{1,2}$ and $\epsilon = \phi_{1,3}$, with scaling dimensions

$$\Delta_\sigma = \frac{1}{2} - \frac{3}{2(m+1)}, \quad \Delta_\epsilon = 2 - \frac{4}{m+1}. \quad (5.1)$$

The central charge is

$$c = 1 - \frac{6}{m(m+1)}. \quad (5.2)$$

From this point on, we solve for m in terms of Δ_σ and use Δ_σ as our independent parameter. For instance, $\Delta_\epsilon = \frac{2}{3}(4\Delta_\sigma + 1)$. The 2d Ising CFT corresponds to $\Delta_\sigma = \frac{1}{8}$.

The operator σ has a null descendant at level 2, so its correlators satisfy a differential equation

$$\left(\mathcal{L}_{-2} - \frac{3}{2(\Delta_\sigma + 1)} \partial_z^2 \right) \langle \sigma(z) \mathcal{O}_1(z_1) \cdots \mathcal{O}_n(z_n) \rangle = 0, \quad (5.3)$$

where

$$\mathcal{L}_{-2} = \sum_{i=1}^n \left(\frac{1}{z - z_i} \partial_{z_i} + \frac{h_i}{(z - z_i)^2} \right), \quad (5.4)$$

and $h_i = \Delta_i/2$ for scalar \mathcal{O}_i . Applying this differential operator to the ansatz

$$\langle \sigma(x_1) \sigma(x_2) \sigma(x_3) \sigma(x_4) \rangle = \frac{1}{x_{12}^{2\Delta_\sigma} x_{34}^{2\Delta_\sigma}} g^{(\Delta_\sigma)}(z, \bar{z}) \quad (5.5)$$

gives a hypergeometric equation that can be solved as

$$g^{(\Delta_\sigma)}(z, \bar{z}) = G_1(z) G_1(\bar{z}) + N(\Delta_\sigma) G_2(z) G_2(\bar{z}), \quad (5.6)$$

$$G_1(z) = (1-z)^{-\Delta_\sigma} {}_2F_1 \left(\frac{1-2\Delta_\sigma}{3}, -2\Delta_\sigma, \frac{2(1-2\Delta_\sigma)}{3}; z \right), \quad (5.7)$$

$$G_2(z) = (1-z)^{\frac{1+\Delta_\sigma}{3}} z^{\frac{1+4\Delta_\sigma}{3}} {}_2F_1 \left(\frac{2(1+\Delta_\sigma)}{3}, 1+2\Delta_\sigma, \frac{4(1+\Delta_\sigma)}{3}; z \right). \quad (5.8)$$

²⁵In the range $\frac{1}{2} \leq \Delta_\sigma \leq 1$, the function $g^{(\Delta_\sigma)}$ may still saturate the bounds on c and Δ_ϵ , but this range has not been sufficiently explored and the conclusive numerical evidence is lacking.

Crossing symmetry then fixes the normalization $N(\Delta_\sigma)$ to be

$$N(\Delta_\sigma) = \frac{2^{1-\frac{8(1+\Delta_\sigma)}{3}} \Gamma\left(\frac{2}{3} - \frac{4\Delta_\sigma}{3}\right)^2 \Gamma(1+2\Delta_\sigma)^2 \left(-\cos\left(\frac{\pi(1+4\Delta_\sigma)}{3}\right) + \sin\left(\frac{\pi(1+16\Delta_\sigma)}{6}\right)\right)}{\pi \Gamma\left(\frac{7}{6} + \frac{2\Delta_\sigma}{3}\right)^2}. \quad (5.9)$$

Notice that $N(\Delta_\sigma)$ is positive for $0 < \Delta_\sigma < 1$.

The interpretation of this solution is that the only Virasoro primaries appearing in the OPE $\sigma \times \sigma$ are $\mathbf{1}$ and ϵ . This is consistent with the OPE structure of the degenerate fields $\phi_{r,s}$. The two terms in the RHS of (5.6) are the holomorphic times anti-holomorphic Virasoro conformal blocks of $\mathbf{1}$ and ϵ , respectively. The function $N(\Delta_\sigma)$ is the square of the $\sigma\sigma\epsilon$ OPE coefficient.

The solution (5.6) can be further decomposed into $SL(2, \mathbb{C})$ conformal blocks, which are the ones used in our numerical analysis. We can then isolate the squared OPE coefficients of the $SL(2, \mathbb{C})$ primaries in the $\sigma \times \sigma$ OPE. The low-lying spectrum of $SL(2, \mathbb{C})$ primaries and their OPE coefficients plotted in Figures 19 and 20 were obtained this way. These states are Virasoro descendants of $\mathbf{1}$ and ϵ . When we do this computation, we lump together all $SL(2, \mathbb{C})$ primaries of the same dimension and spin.

An interesting feature of the solution (5.6) is that its expansion into $SL(2, \mathbb{C})$ conformal blocks has positive coefficients for all $\frac{1}{8} \leq \Delta_\sigma \leq \frac{1}{2}$. This has been checked to a high level in [43], and we checked it to even higher levels. It seems likely that positivity holds for all levels. This may seem surprising because by the non-unitarity theorem [44, 45] the Verma modules of the $\phi_{r,s}$ operators do contain negative norm states unless m is an integer. What happens is that the norms of these negative norm states always remain much smaller in absolute value than the norms of the many positive norm states present at the same level. Intuitively this is to be expected since these negative norms must vanish at the minimal models, so there is not enough room for them to grow. Thus the total contribution per level never becomes negative. We checked explicitly that this is precisely what happens in the interval between the first and the second unitary models. In this interval the first negative norm descendant of $\mathbf{1}(\epsilon)$ occurs at level 12(6), respectively.

The situation changes qualitatively for $\Delta_\sigma < \frac{1}{8}$, where the four-point function ceases to be unitary. The reason is that in this interval the first negative norm descendant of ϵ occurs at level two, and since it's the only descendant on this level there is no room for the cancelation. The operator in question is $\mathcal{O}_2 \equiv (L_{-2} - \frac{3}{2(\Delta_\sigma+1)}L_{-1}^2)\epsilon$ of spin 2 and a norm given by

$$\frac{(1+\Delta_\sigma)(2+5\Delta_\sigma)(8\Delta_\sigma-1)}{6(7+4\Delta_\sigma)(5+8\Delta_\sigma)}. \quad (5.10)$$

We see explicitly that the norm becomes negative when $\Delta_\sigma < \frac{1}{8}$, to the left of the Ising point.

In Figure 20, \mathcal{O}_2 's dependence on Δ_σ is marked by the red squares. The OPE coefficient goes to zero at the Ising point. This dependence is reproduced well by our numerical solution. Another spin 2 state which decouples at the Ising point can be obtained by acting on \mathcal{O}_2 by $(L_{-2} + \dots)(\bar{L}_{-2} + \dots)$, where \dots is fixed to get a quasiprimary. This is the ‘‘question

mark” state which is missed by our numerics. Finally, the decoupling scalar in Figure 19 is $(\bar{L}_{-2} - \frac{3}{2(\Delta_\sigma+1)}\bar{L}_{-1}^2)\mathcal{O}_2$. Its norm thus goes to zero as the square of (5.10).

6 Technical Details

While the computational algorithms used in this paper share many conceptual similarities with our previous work [1], there are two major technical novelties.

Firstly, we switched to a very efficient representation for conformal block derivatives recently developed in [46, 47, 4]. In our previous work [1], we had to discretize Δ with a small step and precompute large tables of conformal blocks and their derivatives corresponding to the discretization. The new representation is sufficiently fast to evaluate conformal blocks inside the simplex algorithm, and allows us to dispense with the discretization.

Secondly, we switched to the “primal” method instead of the “dual” method used in prior work (see section 2.3). Although the two methods are formally equivalent, the primal one has an advantage that at every step of the extremization algorithm we have a valid solution to the crossing symmetry constraint. We implemented our own version of Dantzig’s primal simplex method algorithm, capable of dealing with a continuum of constraints, and using multiple precision arithmetic with $O(100)$ significant digits. Demands of final accuracy and numerical stability render the standard double precision arithmetic (i.e., 16 significant digits) insufficient for our needs (see note 37 for a discussion why).

There are many small subtleties and tweaks which go into the implementation of the above two ideas; they will be described in detail in the rest of this section. It should be noted that the numerical bootstrap field is still rapidly developing and our algorithms will continue to improve dramatically in the foreseeable future. For this reason it doesn’t seem worthwhile yet to release general-purpose code for doing these computations. But we will be happy to provide anyone interested with a current version of our code.

6.1 Partial Fraction Representation for Conformal Blocks

In this section, we describe a representation for conformal blocks that is efficient to compute and well suited for our optimization algorithm. It is based on the series expansion of [46] and the idea of rational approximations introduced in [4].

Conformal blocks are best studied in the radial coordinates of [46]. Via a conformal transformation, we can always place four operators $\sigma(x_1), \dots, \sigma(x_4)$ on a two-plane, so that $x_3 = 1, x_4 = -1$ lie on the unit circle, and $x_1 = -\rho, x_2 = \rho$ lie on a smaller circle around the origin (Figure 21). The complex coordinate $\rho = re^{i\theta}$ is related to the usual conformal cross-ratios via

$$\rho = \frac{z}{(1 + \sqrt{1 - z})^2}. \tag{6.1}$$

A conformal block is then given by inserting all the states in a conformal multiplet (i.e.




Figure 21: Configuration of points for radial quantization in the ρ -coordinate [46].

a primary \mathcal{O} and its descendants $\partial\mathcal{O}, \partial^2\mathcal{O}, \dots$) on a sphere separating x_1, x_2 from x_3, x_4 , in radial quantization,

$$x_{12}^{-2\Delta_\sigma} x_{34}^{-2\Delta_\sigma} G_{\Delta, \ell}(u, v) = \sum_{\alpha=\mathcal{O}, \partial\mathcal{O}, \dots} \frac{\langle 0 | \sigma(x_3) \sigma(x_4) | \alpha \rangle \langle \alpha | \sigma(x_1) \sigma(x_2) | 0 \rangle}{\langle \alpha | \alpha \rangle}. \quad (6.2)$$

By classifying the states α according to their representations under dilatation and rotations, this sum can be written [46]

$$G_{\Delta, \ell}(r, \eta) = \sum_{n=0}^{\infty} \sum_j B_{n, j}(\Delta, \ell) r^{\Delta+n} \frac{j!}{(2\nu)_j} C_j^{(\nu)}(\eta) \quad (6.3)$$

where $r = |\rho|$, $\eta = \cos\theta = \frac{\rho+\bar{\rho}}{2|\rho|}$, $\nu = \frac{d-2}{2}$, and $C_j^{(\nu)}(\eta)$ are Gegenbauer polynomials. The coefficients $B_{n, j}(\Delta, \ell)$ express contributions of descendants of spin j at level n of the multiplet; they are rational functions of the dimension Δ . This follows directly from the expression (6.2) for $G_{\Delta, \ell}$ as a sum over states: each term in the numerator and denominator is a polynomial in Δ that can be computed using the conformal algebra.

For our purposes, we will be interested in computing conformal blocks and their derivatives around the crossing-symmetric point $z = \bar{z} = 1/2$. This corresponds to $r = r_* \equiv 3 - 2\sqrt{2} \approx 0.17$, so the series (6.3) is a rapidly convergent expansion at this point.²⁶ To get a good approximation, we can truncate it at some large but finite value of n , with the result

$$G_{\Delta, \ell}(r, \eta) \approx r^\Delta \frac{P_\ell(r, \eta, \Delta)}{Q_\ell(\Delta)}, \quad (6.4)$$

where P_ℓ, Q_ℓ are polynomials in Δ . Taking derivatives around the crossing-symmetric point, and expanding the resulting rational function of Δ in partial fractions, we can write

$$\partial_r^m \partial_\eta^n G_{\Delta, \ell}(r, \eta) \Big|_{r=r_*, \eta=1} \approx r_*^\Delta \left(p_\ell^{m, n}(\Delta) + \sum_i \frac{a_{\ell, i}^{m, n}}{\Delta - \Delta_i} \right), \quad (6.5)$$

²⁶For the conformal blocks corresponding to identical external scalars used here, the expansion is actually in powers of r^2 [46].

where $p_\ell^{m,n}(\Delta)$ are polynomials and $a_{\ell,i}^{m,n}$ are numerical coefficients. As Table 4 below shows, there are only simple poles in Δ except if $d = 2, 4, 6, \dots$ when there are also double poles. Conformal blocks vary continuously in d , so the double poles get resolved into a pair of simple poles when d is slightly perturbed away from an even integer.²⁷

This representation of the conformal blocks in terms of partial fractions has the virtue that once the data $p_\ell^{m,n}$ and $a_{\ell,i}^{m,n}$ have been computed, we can calculate the blocks at any value of the dimension Δ extremely rapidly and with very high precision. This will be crucial in our optimization algorithm. Now, having established what representation of the conformal blocks we would like to use, let us describe two methods for computing it. For this work, we have implemented both methods.

6.1.1 The Casimir Equation

The conformal block $G_{\Delta,\ell}(r, \eta)$ is an eigenvector of the quadratic Casimir of the conformal group. This implies a differential equation for $G_{\Delta,\ell}(r, \eta)$ which can be solved iteratively order by order in r [46]. Using the Casimir equation, one can recursively compute the coefficients $B_{n,j}(\Delta, \ell)$ starting from the initial conditions:

$$B_{0,j}(\Delta, \ell) = 4^\Delta \delta_{j\ell}, \quad (6.6)$$

expressing the fact that the primary itself is the only state on level zero of the multiplet. The 4^Δ fixes the normalization of conformal blocks to be the same as in [1]. Once $B_{n,j}(\Delta, \ell)$ are known, it is straightforward to obtain $p_\ell^{m,n}, a_{\ell,i}^{m,n}$.

In this work, we did not literally use this method but a slight variation due to [47]. The point is that Eq. (6.3) contains more information than needed: it can be used to recover conformal blocks for any η while in the applications here we only need $\eta \approx 1$. The variation described below focuses directly on $\eta \approx 1$ without having to deal with Gegenbauer polynomials.

One considers first the conformal block restricted to the ‘‘diagonal’’ $\rho = \bar{\rho}$ (i.e. $\eta = 1$):

$$G_{\Delta,\ell}(r) \equiv G_{\Delta,\ell}(r, \eta = 1) = \rho^\Delta \sum_{n=0}^{\infty} b_n(\Delta, \ell) r^n, \quad b_n(\Delta, \ell) = \sum_j B_{n,j}(\Delta, \ell). \quad (6.7)$$

It was shown in [47] that $G_{\Delta,\ell}(r)$ satisfies a fourth-order ordinary differential equation (which becomes third-order for $\ell = 0$). Using this equation, we evaluate $b_n(\Delta, \ell)$ to a very high order ($n = 120$ was used in most computations, and a few results were checked at $n = 200$) starting from $b_0(\Delta, \ell) = 4^\Delta$. The equation has a regular singular point at $\rho = 0$, which is why we can evaluate the whole series expansion starting from just one initial condition for the leading term. Knowing $b_n(\Delta, \ell)$, we obtain coefficients $p_\ell^{m,0}, a_{\ell,i}^{m,0}$ in (6.5).

²⁷We used this feature in our $d = 2$ computations presented in section 5. Instead of dealing with double poles, we used our generic simple-pole code and ran it at $d = 2 + 10^{-5}$. We checked that using $d = 2 + 10^{-7}$ or even $d = 2 + 10^{-15}$ does not change the results.

n_i	Δ_i	ℓ_i	c_i	
$2k$	$1 - \ell - 2k$	$\ell + 2k$	$c_1(k)$	$k = 1, 2, \dots$
$2k$	$1 + \nu - k$	ℓ	$c_2(k)$	$k = 1, 2, \dots$
$2k$	$1 + \ell + 2\nu - 2k$	$\ell - 2k$	$c_3(k)$	$k = 1, 2, \dots, \lfloor \ell/2 \rfloor$

Table 4: The positions of poles of $G_{\Delta,\ell}$ in Δ and their associated data. There are three types of poles, corresponding to the three rows in the table. The first two types exist for all positive integer k , while the third type exists for positive integer $k \leq \lfloor \ell/2 \rfloor$. The coefficients $c_1(k), c_2(k), c_3(k)$ are given in Eqs. (6.11).

To obtain derivatives normal to the diagonal, we then use the quadratic Casimir partial differential equation, solving it *à la* Cauchy-Kovalevskaya in a power-series expansion around the diagonal. This idea was already used in [1], where it was shown that all normal and mixed derivatives of conformal blocks could be reduced in this way to derivatives along the diagonal.

To be precise, we switch halfway to the variables a and b defined as in [1],

$$z = (a + \sqrt{b})/2, \quad \bar{z} = (a - \sqrt{b})/2, \quad (6.8)$$

so that $b = 0$ is the diagonal, and the crossing symmetric point $z = \bar{z} = 1/2$ corresponds to $a = 1, b = 0$. Since the conformal blocks are symmetric under $z \leftrightarrow \bar{z}$, the expansion around $b = 0$ contains integer powers of b . So, we first evaluate the diagonal derivatives $\partial_r^m G_{\Delta,\ell}$ as described above, then do a change of variables $\rho \rightarrow a$ to express the diagonal derivatives $\partial_a^m G_{\Delta,\ell}|_{a=1}$, and finally use the Cauchy-Kovalevskaya method, precisely as described in section 4 of [1], to obtain the mixed and normal derivatives $\partial_a^m \partial_b^n G_{\Delta,\ell}|_{a=1,b=0}$. For all these derivatives we get a representation of the form (6.5).

6.1.2 Recursion Relations

An alternate way to compute the partial-fraction representation (6.5) was developed in [4]. Here, we review this method and include additional details about its implementation. The idea is to use a recursion relation expressing conformal blocks as a sum over poles in Δ , where the residue at each pole is itself a conformal block:

$$\begin{aligned} h_{\Delta,\ell}(r, \eta) &\equiv r^{-\Delta} G_{\Delta,\ell}(r, \eta) \\ h_{\Delta,\ell}(r, \eta) &= h_\ell^{(\infty)}(r, \eta) + \sum_i \frac{c_i r^{n_i}}{\Delta - \Delta_i} h_{\Delta_i + n_i, \ell_i}(r, \eta), \end{aligned} \quad (6.9)$$

with

$$h_\ell^{(\infty)}(r, \eta) = \frac{\ell!}{(2\nu)_\ell} \frac{C_\ell^{(\nu)}(\eta)}{(1-r^2)^\nu \sqrt{(1+r^2)^2 - 4r^2\eta^2}}. \quad (6.10)$$

The Δ_i above are special values of the dimension where a descendant state $|\alpha\rangle$ in (6.2) can become null. By definition, these degenerate values are always below the unitarity

bound (2.3). The null descendant $|\alpha\rangle$ has dimension $\Delta_i + n_i$ and spin ℓ_i , and each pole comes with a numerical factor c_i . For the reader's convenience, we summarize this data in Table 4 and Eq. (6.11):

$$\begin{aligned}
c_1(k) &= -\frac{k(2k)!^2 (\ell + 2\nu)_{2k}}{2^{4k-1} k!^4 (\ell + \nu)_{2k}}, \\
c_2(k) &= -\frac{k(\nu + \ell - k)(\nu)_k(1 - \nu)_k \left(\frac{\nu + \ell + 1 - k}{2}\right)_k^2}{k!^2(\nu + \ell + k) \left(\frac{\nu + \ell - k}{2}\right)_k^2}, \\
c_3(k) &= -\frac{k(2k)!^2 (1 + \ell - 2k)_{2k}}{2^{4k-1} k!^4 (1 + \nu + \ell - 2k)_{2k}}.
\end{aligned} \tag{6.11}$$

Note that each term in the sum over poles Δ_i is suppressed by at least r^2 . Thus, by iterating the recursion relation (6.9), starting with the initial term $h_\ell^{(\infty)}$, we rapidly converge to the correct value of the conformal block. In practice, it is convenient to take derivatives first and perform this iteration numerically at a given point $(r, \eta) = (r_*, 1)$, while keeping Δ as a variable. Specifically, let us define the vector of derivatives

$$\begin{aligned}
\mathbf{h}_\ell(\Delta) &= [\partial_r^0 \partial_\eta^0 h_{\Delta, \ell}(r_*, 1) \quad \partial_r^1 \partial_\eta^0 h_{\Delta, \ell}(r_*, 1) \quad \cdots]^T \\
&= \mathbf{h}_\ell^{(\infty)} + \sum_i \frac{\mathbf{d}_{\ell, i}}{\Delta - \Delta_i},
\end{aligned} \tag{6.12}$$

where we include all derivatives $\partial_r^m \partial_\eta^n$ with $m + n \leq 2K$ for some K . The vector $\mathbf{h}_\ell^{(\infty)}$ is given simply by derivatives of the known function (6.10), which can be computed beforehand. Multiplication by r is represented on the space of derivatives by a matrix \mathbf{R} , so Eq. (6.9) implies

$$\mathbf{d}_{\ell, i} = c_i(\ell) \mathbf{R}^{n_i} \left(\mathbf{h}_\ell^{(\infty)} + \sum_j \frac{\mathbf{d}_{\ell, j}}{\Delta_i(\ell) + n_i - \Delta_j(\ell_i)} \right). \tag{6.13}$$

Iterating this equation numerically, we can compute the residues $\mathbf{d}_{\ell, i}$. The number of iterations, together with the number of poles in the ansatz (6.12) is roughly equivalent to the order at which we truncate the Gegenbauer expansion (6.3). In this work, we keep poles up to roughly $k = 100$ and perform 100 iterations.

Finally, to recover the vector of derivatives for $G_{\Delta, \ell}$ itself, we should multiply by the matrix \mathbf{R}^Δ . This matrix takes the form $\mathbf{R}^\Delta = r^\Delta \mathbf{S}(\Delta)$, where $\mathbf{S}(\Delta)$ is a matrix polynomial in Δ . Decomposing $\mathbf{S}(\Delta) \mathbf{h}_\ell(\Delta)$ into partial fractions then yields the representation (6.5). For example, the coefficients $\mathbf{a}_{\ell, i} = (a_{\ell, i}^{m, n})$ in section 6.1 are given by

$$\mathbf{a}_{\ell, i} = \text{Res}_{\Delta \rightarrow \Delta_i} \mathbf{S}(\Delta) \frac{\mathbf{d}_{\ell, i}}{\Delta - \Delta_i} = \mathbf{S}(\Delta_i) \mathbf{d}_{\ell, i}. \tag{6.14}$$

6.2 A Customized Simplex Method

In this section, we describe our procedure for optimizing over $\mathcal{C}_{\Delta, \sigma}$. The underlying algorithm is the Simplex Method, due to Dantzig. We will first review this algorithm, and then discuss its specialization to our case of interest.

6.2.1 The Primal Simplex Method

The material in this section is standard in the mathematics and computer science literature. We include it to establish notation and because it may be unfamiliar to physicist readers. We will essentially follow the presentation of [48].

Given vectors $\mathbf{c} \in \mathbb{R}^n$, $\mathbf{b} \in \mathbb{R}^m$ and a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, we would like to minimize the *objective function*

$$\mathbf{c} \cdot \mathbf{x} = c_1 x_1 + \cdots + c_n x_n$$

over $\mathbf{x} \in \mathbb{R}^n$ such that

- $x_i \geq 0$, and
- $\mathbf{A}\mathbf{x} = \mathbf{a}_1 x_1 + \cdots + \mathbf{a}_n x_n = \mathbf{b}$.

We assume $n > m$, so that the space of possible \mathbf{x} 's has positive dimension.

The space of possible \mathbf{x} 's is a convex polytope. Because of convexity, the minimum we seek is always realized at a vertex of this polytope (though it may be non-unique). So it suffices to minimize $\mathbf{c} \cdot \mathbf{x}$ over vertices.

At a vertex, as many as possible of the inequalities $x_i \geq 0$ are saturated. Since \mathbf{x} lives in n -dimensions and is subject to m equality constraints $\mathbf{A}\mathbf{x} = \mathbf{b}$, generically $n - m$ inequalities can be saturated. This leaves m nonzero coordinates x_{j_1}, \dots, x_{j_m} .²⁸ The equality $\mathbf{A}\mathbf{x} = \mathbf{b}$ then expresses \mathbf{b} as a nonnegative sum of the corresponding vectors $\mathbf{a}_{j_1}, \dots, \mathbf{a}_{j_m}$, which are called *basic vectors*,

$$\mathbf{a}_{j_1} x_{j_1} + \cdots + \mathbf{a}_{j_m} x_{j_m} = \mathbf{b}. \quad (6.15)$$

The remaining \mathbf{a}_i with $x_i = 0$ are called *nonbasic*. Specifying a vertex is equivalent to specifying a set of m basic vectors.

The idea of the simplex algorithm is to travel from vertex to vertex along the polytope edges, following the direction of steepest descent. Let's assume that we have found a vertex of our polytope (we address the question of finding an initial vertex later), and describe how to pass to the next vertex.

Suppose our starting vertex is characterized by m basic vectors $\mathbf{a}_{j_1}, \dots, \mathbf{a}_{j_m}$ with nonzero coordinates x_{j_1}, \dots, x_{j_m} . For convenience, let us partition \mathbf{A} into an $m \times m$ matrix \mathbf{A}_B whose columns are basic vectors, and an $m \times (n - m)$ matrix \mathbf{A}_N whose columns are the remaining nonbasic vectors,

$$\mathbf{A} = [\mathbf{A}_B \mid \mathbf{A}_N] \quad (6.16)$$

We similarly partition \mathbf{x} into basic coordinates $\mathbf{x}_B = (x_{j_1}, \dots, x_{j_m})^T$ and the remaining nonbasic coordinates \mathbf{x}_N (all of which vanish at our vertex).

²⁸At non-generic vertices, it's possible that extra inequalities can be saturated, so that some of the x_{j_i} actually vanish.

Now consider adjusting some nonbasic coordinate x_k away from zero. To preserve $\mathbf{Ax} = \mathbf{b}$, we must simultaneously adjust \mathbf{x}_B . We have

$$\mathbf{A}_B \mathbf{x}_B + \mathbf{a}_k x_k = \mathbf{b} \quad \Longrightarrow \quad \mathbf{x}_B = \mathbf{A}_B^{-1}(\mathbf{b} - \mathbf{a}_k x_k), \quad (6.17)$$

so that our objective function becomes

$$\mathbf{c}_B \cdot \mathbf{x}_B + c_k x_k = \mathbf{c}_B^T \mathbf{A}_B^{-1} \mathbf{b} + (c_k - \mathbf{c}_B^T \mathbf{A}_B^{-1} \mathbf{a}_k) x_k. \quad (6.18)$$

To decrease $\mathbf{c} \cdot \mathbf{x}$ as quickly as possible, we should choose k that minimizes the quantity in parentheses, known as the *reduced cost*,

$$\begin{aligned} k_* &= \operatorname{argmin}_k \operatorname{RC}_k, \\ \operatorname{RC}_k &\equiv (c_k - \mathbf{c}_B^T \mathbf{A}_B^{-1} \mathbf{a}_k). \end{aligned} \quad (6.19)$$

If the minimum reduced cost is nonnegative, then our starting vertex is already a minimum, and the algorithm terminates. Assume instead that the minimum reduced cost is negative. Once we've picked k_* , we turn on x_{k_*} as much as possible until one of the original basic coordinates goes to zero,

$$x_\ell = (\mathbf{A}_B^{-1} \mathbf{b} - \mathbf{A}_B^{-1} \mathbf{a}_{k_*} x_{k_*})_\ell = 0. \quad (6.20)$$

We choose ℓ to be the first index for which this happens (this is commonly known as the “ratio test”). When x_ℓ becomes zero, the result is a new set of m basic vectors where \mathbf{a}_ℓ has been replaced by \mathbf{a}_{k_*} . This defines a new vertex with strictly smaller objective function $\mathbf{c} \cdot \mathbf{x}$. By repeating this process, we eventually reach a minimum.

6.2.2 Choosing an Initial Vertex

To run the iteration described above, we need an initial vertex. We can find one by solving an auxiliary optimization problem (typically called “phase 1” of the algorithm, while the main optimization stage described above is called “phase 2”). We extend \mathbf{x} with m additional *slack variables* $\mathbf{s} = (s_1, \dots, s_m)^T$,

$$\mathbf{x}' = \left(\begin{array}{c} \mathbf{s} \\ \mathbf{x} \end{array} \right), \quad (6.21)$$

and extend \mathbf{A} with *slack vectors*,

$$\mathbf{A}' = \left(\begin{array}{cccc|c} b_1 & 0 & \dots & 0 & \mathbf{A} \\ 0 & b_2 & \dots & 0 & \\ \vdots & \vdots & \ddots & \vdots & \\ 0 & 0 & \dots & b_m & \end{array} \right).$$

These are designed so that we can trivially satisfy the conditions

$$\mathbf{A}' \mathbf{x}' = \mathbf{b} \quad \text{and} \quad x'_i \geq 0 \quad (i = 1, \dots, n + m) \quad (6.22)$$

by choosing the vertex $s_i = 1$, $\mathbf{x} = 0$. Starting from this initial condition, we can now try to minimize

$$\mathbf{c}' \cdot \mathbf{x}' \equiv \sum_{i=1}^m s_i \quad (6.23)$$

using the algorithm described above. This objective function is bounded below because the s_i 's are nonnegative. If the minimum is zero, then every s_i vanishes and we've found a good vertex for starting our original optimization problem. If the minimum is nonzero, then no such vertex exists and the original optimization problem is infeasible.

Hot Start

This is a slightly more involved strategy for an initial vertex search, which may give a considerable speed up, provided that we have a good guess for an initial basis. Let this guess be given by an $n \times n$ "trial basis" submatrix $\tilde{\mathbf{A}}_B$ of \mathbf{A} . The trial coordinates are given by

$$\tilde{\mathbf{x}}_B = \tilde{\mathbf{A}}_B^{-1} \mathbf{b}. \quad (6.24)$$

If all of them are positive, then our guess was perfect, and we can go straight to phase 2. Otherwise, we extend the matrix \mathbf{A} by the *negatives* of all columns in $\tilde{\mathbf{A}}_B$ whose trial coordinates are negative. If the initial guess was reasonably good, there will be only a few such vectors. The added vectors plus all the remaining vectors of $\tilde{\mathbf{A}}_B$ form a good basis for the extended problem. We then form the objective function given by the sum of the coordinates of all the added vectors, and attempt to minimize it to zero.

6.2.3 Adaptation to Optimization over $\mathcal{C}_{\Delta_\sigma}$

Let us describe the space $\mathcal{C}_{\Delta_\sigma}$ in the language of the previous sections. Crossing symmetry Eq. (2.4) can be written

$$\sum_{\Delta, \ell} p_{\Delta, \ell} F_{\Delta, \ell}^{\Delta_\sigma}(u, v) = -F_{0,0}^{\Delta_\sigma}(u, v), \quad (6.25)$$

where

$$F_{\Delta, \ell}^{\Delta_\sigma}(u, v) \equiv u^{-\Delta_\sigma} G_{\Delta, \ell}(u, v) - v^{-\Delta_\sigma} G_{\Delta, \ell}(v, u) \quad (6.26)$$

and $F_{0,0}^{\Delta_\sigma}$ corresponds to the unit operator.

As explained in section 2.3, in practice we work with the spaces $\mathcal{C}_{\Delta_\sigma}^{(N)}$, given by truncating the crossing symmetry constraint to N derivatives around the crossing-symmetric point $z = \bar{z} = 1/2$. Let us define vectors of derivatives $\mathbf{F}_{\Delta, \ell}^{\Delta_\sigma}$ with components

$$(\mathbf{F}_{\Delta, \ell}^{\Delta_\sigma})^{(m, n)} = \partial_z^m \partial_{\bar{z}}^n F_{\Delta, \ell}^{\Delta_\sigma}(u, v) \Big|_{z=\bar{z}=1/2}, \quad (6.27)$$

where $u = z\bar{z}$, $v = (1-z)(1-\bar{z})$ as usual. Because $F_{\Delta, \ell}^{\Delta_\sigma}(u, v)$ is antisymmetric under $u \leftrightarrow v$, these derivatives are nonzero only if $m + n$ is odd. Further, since u and v are invariant

under $z \leftrightarrow \bar{z}$, it suffices to take $m \geq n$. Demanding $m + n \leq 2K$ for integer K , we have $N = \frac{K(K+1)}{2}$ nontrivial components for $\mathbf{F}_{\Delta,\ell}^{\Delta\sigma}$. This is the truncation parameter N used to label plots throughout the paper. For example $N = 231$ corresponds to $K = 21$.

The vectors $\mathbf{F}_{\Delta,\ell}^{\Delta\sigma}$ are related by a linear transformation to the derivatives of conformal blocks discussed in section 6.1. We first transform from derivatives in r, η coordinates (section 6.1.2) to z, \bar{z} coordinates, and then multiply by an additional matrix to account for multiplication by $u^{-\Delta\sigma}$ and $v^{-\Delta\sigma}$ in Eq. (6.26).

Alternatively, we can work with the a, b coordinates from section 6.1.1 and define

$$(\mathbf{F}_{\Delta,\ell}^{\Delta\sigma})^{(m,n)} = \partial_a^m \partial_b^n F_{\Delta,\ell}^{\Delta\sigma}(u, v) \Big|_{a=1, b=0}, \quad (6.28)$$

where $m, n \geq 0$ arbitrary, m is odd (otherwise the derivative vanishes), and $m + 2n \leq 2K - 1$. Working with these derivatives or with (6.27) corresponds to a different choice of basis in the same N -dimensional space and gives the same constraints.

After taking derivatives, the crossing equation takes the familiar form

$$\sum_i \mathbf{a}_i x_i = \mathbf{b}, \quad x_i \geq 0, \quad (6.29)$$

with the identifications

$$\begin{aligned} i &\leftrightarrow (\Delta, \ell), \\ x_i &\leftrightarrow p_{\Delta,\ell}, \\ \mathbf{a}_i &\leftrightarrow \mathbf{F}_{\Delta,\ell}^{\Delta\sigma}, \\ \mathbf{b} &\leftrightarrow -\mathbf{F}_{0,0}^{\Delta\sigma}. \end{aligned} \quad (6.30)$$

p_T -maximization now resembles the linear program studied in the previous section, with objective function $\mathbf{c} \cdot \mathbf{x} = -p_{d,2}$.

The key difference between optimization over $\mathcal{C}_{\Delta\sigma}^{(N)}$ and a typical linear program is that the number of vectors $\mathbf{F}_{\Delta,\ell}^{\Delta\sigma}$ is continuously infinite, whereas before we only considered a finite collection of vectors \mathbf{a}_i . Such optimization problems are called a *semi-infinite programs*. In the simplex algorithm, finding the minimum reduced cost Eq. (6.19) now requires minimizing over infinitely many vectors.

For concreteness, consider a vertex with N basic vectors

$$\mathbf{F}_{\Delta_1,\ell_1}^{\Delta\sigma}, \dots, \mathbf{F}_{\Delta_N,\ell_N}^{\Delta\sigma}. \quad (6.31)$$

We call the dimensions and spins $\{(\Delta_1, \ell_1), \dots, (\Delta_N, \ell_N)\}$ the ‘‘spectrum’’ at this vertex. When we turn on a nonbasic coefficient $p_{\Delta,\ell}$, the reduced cost is now a function of Δ and ℓ ,

$$\text{RC}_{\Delta,\ell} = c_{\Delta,\ell} - \mathbf{c}_B^T \mathbf{A}_B^{-1} \mathbf{F}_{\Delta,\ell}^{\Delta\sigma}, \quad (6.32)$$

where

$$\mathbf{A}_B = (\mathbf{F}_{\Delta_1,\ell_1}^{\Delta\sigma} \ \cdots \ \mathbf{F}_{\Delta_N,\ell_N}^{\Delta\sigma}), \quad (6.33)$$

$$\mathbf{c}_B = (c_{\Delta_1,\ell_1} \ \cdots \ c_{\Delta_N,\ell_N})^T, \quad (6.34)$$

and $c_{\Delta,\ell}$ is the coefficient of $p_{\Delta,\ell}$ in the objective function.

Suppose $c_{\Delta,\ell}$ is nonzero only for finitely many (Δ, ℓ) . For example, in p_T -maximization, we have $c_{d,2} = -1$, while $c_{\Delta,\ell} = 0$ otherwise. If there are any such (Δ, ℓ) which are not already in the basis, we can simply compute the reduced cost for these special (Δ, ℓ) . This is not even needed in p_T -maximization, since the stress tensor will always be in the basis.²⁹ The challenge is scanning over the infinite (Δ, ℓ) for which $c_{\Delta,\ell}$ vanishes. In this case, we must minimize

$$\text{RC}_{\Delta,\ell} = -\mathbf{c}_B^T \mathbf{A}_B^{-1} \mathbf{F}_{\Delta,\ell}^{\Delta\sigma} \quad (6.35)$$

over all dimensions and spins (Δ, ℓ) which can appear in the spectrum.

By the results of section 6.1, $\text{RC}_{\Delta,\ell}$ can be arbitrarily well approximated by a partial fraction expansion in Δ ,

$$\text{RC}_{\Delta,\ell} \approx r_*^\Delta \left(q_\ell(\Delta) + \sum_i \frac{s_{\ell,i}}{\Delta - \Delta_i} \right), \quad (6.36)$$

where $r_* = 3 - 2\sqrt{2} \approx 0.17$. The polynomials q_ℓ and residues $s_{\ell,i}$ are given by dot products of $-\mathbf{c}_B^T \mathbf{A}_B^{-1}$ with the vectors $\mathbf{p}_\ell(\Delta)$ and $\mathbf{a}_{\ell,i}$ appearing in Eq. (6.5). This expression for $\text{RC}_{\Delta,\ell}$ has several nice properties. Firstly, since $r_* \ll 1$, it's clear that $|\text{RC}_{\Delta,\ell}|$ falls off quickly with Δ so that it suffices to minimize over a finite range of dimensions Δ .³⁰ Secondly, we can evaluate it efficiently on a computer for any Δ .

Our modification of the simplex algorithm is to approximate the reduced cost by Eq. (6.36) and minimize over Δ using an efficient univariate minimization algorithm, described in the next section. In initial bootstrap studies (using out-of-the-box solvers), the strategy was to discretize the allowed dimensions Δ [8]. This meant minimizing the reduced cost by scanning over every dimension in this discrete set—an expensive operation. Precision errors could also be introduced by the discretization. Our new algorithm evades these difficulties.³¹

6.2.4 A Strategy for Reduced Cost Minimization

To minimize the reduced cost, we must minimize functions of the form (6.36) over dimensions Δ satisfying the unitarity bound (2.3). We will restrict to $\ell \leq \ell_{\max}$ for some large ℓ_{\max}

²⁹Except possibly at one intermediate step after the hot start, see section 6.2.5.

³⁰The simplex algorithm can proceed as long as we always find a negative reduced cost at every step. Thus we can rescale $\text{RC}_{\Delta,\ell}$ by any positive function of Δ , and we will still eventually find the correct optimum. However, different rescalings cause the search to proceed in different ways. For example, if we strip off the factor r_*^Δ from $\text{RC}_{\Delta,\ell}$, the simplex algorithm will proceed along a different path, favoring larger values of Δ in the intermediate steps. By the exponential decoupling theorem of [9], low-lying operators will dominate the constraints of crossing symmetry, and it is more practical to start the search by exploring the low dimensions first. We have found that the normalization in Eq. (6.36) realizes this requirement in practice, and leads to the fastest solution times. This is because this normalization is natural from the point of view of the OPE convergence estimates of [9].

³¹An alternative way to avoid discretizing Δ is to use semidefinite programming [12, 4], see Appendix A.

(typically 40 or 50).³² Further, since $\text{RC}_{\Delta,\ell}$ falls off exponentially quickly with Δ , it suffices to minimize over Δ in some finite interval $[\Delta_{\text{unitarity}}, \Delta_{\text{max}}]$. We typically take $\Delta_{\text{max}} = 50$ or the sliding cutoff $\Delta_{\text{max}} = 50 + \ell$. We checked that our results are insensitive to varying these cutoffs.

Thus, we have the problem of minimizing a smooth univariate function $f(\Delta)$ over an interval $\Delta \in [a, b]$. There are many possible strategies. Here, we present one that works well in practice, though we sacrifice some rigor for speed. This is because the simplex algorithm can proceed as long as we find *some* negative minimum at each step. In other words, it's not absolutely necessary to find the true minimum reduced cost every time. However, we benefit if our algorithm manages to find the true minimum most of the time.³³

Our strategy proceeds in four steps:

1. Recursively divide $[a, b]$ into smaller intervals where the shape of f' is well understood.
2. Find which of these small intervals could contain a local minimum of f .
3. Compute the local minima using Newton's method.
4. Pick the overall minimum among the local minima and the endpoints a, b .

Let us describe these steps in detail. For an interval $[x, y]$, define a quadratic fit for f' around the midpoint $z = \frac{x+y}{2}$,

$$f'_{\text{fit}}(t) \equiv f'(z) + f''(z)(t - z) + \frac{1}{2}f'''(z)(t - z)^2. \quad (6.37)$$

Call the interval $[x, y]$ “good” if this fit agrees closely with the true values at the endpoints,

$$\frac{|f'(x) - f'_{\text{fit}}(x)|}{|f'(x)| + |f'_{\text{fit}}(x)|} \leq \epsilon \quad \text{and} \quad \frac{|f'(y) - f'_{\text{fit}}(y)|}{|f'(y)| + |f'_{\text{fit}}(y)|} \leq \epsilon, \quad (6.38)$$

where ϵ is a small parameter (typically 0.05).

In step 1, we check whether the interval $[a, b]$ is good. If not, we split it into two intervals $[a, \frac{a+b}{2}]$ and $[\frac{a+b}{2}, b]$ and recursively check and split each interval. We stop when we've completely partitioned $[a, b]$ into good intervals.³⁴

In step 2, the good intervals which can contain a local minimum are those where f' is negative on the left endpoint and positive on the right endpoint. This criterion is sufficient

³²This truncation of spins is an approximation. By the unitarity bound, operators with large spin have large dimension, and one can show that the contribution of large dimension operators to a four-point function is exponentially suppressed at $z = \bar{z} = 1/2$ [9]. However, if we investigate the constraints of crossing symmetry near the lightcone ($z \rightarrow 0$ with \bar{z} fixed), then it is known that large spin operators play an important role [18, 19]. It will be interesting to explore this regime in future bootstrap studies.

³³When the simplex algorithm terminates, we do want to be sure of the true minimum. In this case, we can either apply a more rigorous minimization strategy or simply decrease the tolerance parameter ϵ in Eq. (6.38).

³⁴Each time we split an interval, we can re-use our calculation of f' at the endpoints and midpoint in the next level of recursion.

but not necessary for the existence of a zero of f' (further, it's possible that a single interval could contain multiple zeros). This is where we sacrifice rigor for speed. Most computation time is spent evaluating f and its derivatives. Thus, it is fruitful to reduce the number of evaluations even if that requires making assumptions about the shape of f' .

Once we have the zeros of f' isolated into intervals, we can use Newton's method to compute them with high precision (step 3).³⁵ This is the least computationally intensive part of the algorithm, and it is easy to compute within an error of 10^{-30} or smaller.

6.2.5 Hot Start from the Mean Field Theory Spectrum

In section 6.2.2 we described a general hot start strategy for the initial vertex search. This strategy works particularly well in our problem of p_T -maximization, because a good initial guess for a spectrum can be obtained from Mean Field Theory (MFT). Recall that in MFT the $\sigma \times \sigma$ OPE contains operators

$$\sigma(\partial\partial \dots \partial)_{\text{spin } \ell} (\partial^2)^n \sigma \tag{6.39}$$

of even spin ℓ and dimension $\Delta = 2\Delta_\sigma + \ell + 2n$. A trial spectrum for our problem can be obtained by truncating this infinite spectrum to N operators. We found that the best truncation strategy is to keep the N lowest dimension operators (preferring low twist operators whenever the dimension is equal). This strategy gives the smallest number of auxiliary variables and the fastest solution time.

It may happen that the basis obtained by hotstarting from the MFT spectrum does not contain the stress tensor (since the MFT spectrum does not have it). Then one intermediate step is required after hot start and before the main stage of the algorithm, to bring the stress tensor into the basis.

Without hot start, the initial vertex search takes about the same time as the subsequent p_T -maximization. Hot start from the MFT spectrum speeds up the initial vertex search by an order of magnitude or more, thus reducing the total solution time by about a factor 2.

6.2.6 Linear Functionals

In this section we clarify the relationship between the primal simplex method described above and the linear functionals discussed in previous bootstrap studies. At each stage in the simplex algorithm, the reduced cost can be written

$$\text{RC}_{\Delta,\ell} = -\mathbf{c}_B^T \mathbf{A}_B^{-1} \mathbf{F}_{\Delta,\ell}^{\Delta_\sigma} \equiv \Lambda[F_{\Delta,\ell}^{\Delta_\sigma}], \tag{6.40}$$

where Λ is a linear functional. In the case of p_T -maximization, all components of \mathbf{c}_B vanish except for $c_{d,2} = -1$. Further, the columns of \mathbf{A}_B are precisely the basic vectors $\mathbf{F}_{\Delta_i,\ell_i}^{\Delta_\sigma}$.

³⁵In fact, we use a hybrid of Newton's method and binary search which is guaranteed to stay within the interval and find a zero.

Acting with Λ on these basic vectors, we find

$$\Lambda[F_{\Delta_i, \ell_i}^{\Delta\sigma}] = -\mathbf{c}_B^T \mathbf{A}_B^{-1} \mathbf{F}_{\Delta_i, \ell_i}^{\Delta\sigma} = -c_{\Delta_i, \ell_i} = \begin{cases} 1 & \text{if } (\Delta_i, \ell_i) = (d, 2), \\ 0 & \text{otherwise} \end{cases} \quad (6.41)$$

In other words, $\Lambda[F_{\Delta, \ell}^{\Delta\sigma}]$ has a zero when (Δ, ℓ) is in the spectrum (excluding the stress-tensor).

If $\Lambda[F_{\Delta, \ell}^{\Delta\sigma}]$ has a negative minimum, then the simplex algorithm instructs us to swap the corresponding vector into the spectrum, raising the functional up to zero there (Figure 22). By repeating this process, we “push up” all the minima of $\Lambda[F_{\Delta, \ell}^{\Delta\sigma}]$, until we obtain a nonnegative functional Λ_* and the algorithm terminates.³⁶




Figure 22: The change in the functional $\Lambda \rightarrow \Lambda'$ after a single step of the simplex algorithm. We first minimize $\Lambda[F_{\Delta, \ell}^{\Delta\sigma}]$ to find $\Delta_{\min}, \ell_{\min}$. Swapping the corresponding vector into the basis results in a new functional with a zero at $\Delta_{\min}, \ell_{\min}$. After many steps, the minima move up to zero and the functional becomes nonnegative.

The zeros of Λ_* are precisely the dimensions and spins present in the optimal spectrum. This establishes the equivalence of the primal method and the dual (or “extremal functional”) method discussed in [10, 2]. In each step of the primal method, we have a solution to crossing symmetry and a functional of indefinite sign. In each step of the dual method, we have a nonnegative functional but no solution to crossing. Both methods terminate with a solution to crossing and a nonnegative functional.

Nonnegativity, together with continuity in Δ , implies that $\Lambda_*[F_{\Delta, \ell}^{\Delta\sigma}]$ actually has a *double zero* at generic (Δ, ℓ) in the optimal spectrum. (Λ_* can have a single zero if it occurs at the

³⁶The algorithm actually does not terminate after a finite number of steps but instead converges quickly to a solution. One can terminate the algorithm by hand when the minimum reduced cost is sufficiently close to zero, $\min(\text{RC}_{\Delta, \ell}) \geq -\delta$ for small δ . Some of our plots (Figures 1 (right), 4 (right), 5, 10, and 14) were produced with $\delta = 10^{-60}$. We have found that a much less conservative criterion $\delta \lesssim 10^{-12}$ suffices to reach the optimal spectrum with reasonable precision. Another criterion is to terminate if p_T is reduced by less than δ in the last M iterations (we used $M = 1000$, $\delta = 10^{-15}$). The point is that the minimum reduced cost is not a perfect predictor for the actual reduction in p_T , and it may also fluctuate significantly from one step to the other. Thus looking at the change in p_T integrated over many iterations may give a better idea about the progress of the algorithm.

end of an interval of allowed Δ , for instance at the unitarity bound.) At intermediate stages in the simplex algorithm, this manifests as pairs of dimensions that approach each other as we converge to the optimal solution.³⁷ The functional dips negative between each pair and is forced up to zero as they squeeze together. Thus, although our basis at each stage contains N vectors, we end up with between $N/2$ and N operators in the optimal spectrum.

At this point, let us address the claim in section 2.2 that $\mathcal{C}_{\Delta_\sigma}$ is infinite dimensional. For each N we explored, the space of possible spectra is infinite dimensional. Even if we consider only extremal spectra, we have an N -dimensional space of possible objective functions to optimize, and for each objective function we generically observe a different spectrum of dimensions Δ_i . Together, these optimal four-point functions span an infinite dimensional space. It would be extremely surprising if this phenomenon ceased to be true at sufficiently large N .

6.2.7 Δ_ϵ -Maximization

To optimize a nonlinear objective function like Δ_ϵ , we must combine the simplex algorithm with an additional search. For example, in Δ_ϵ -maximization we start by assuming all scalars have dimension $\Delta \geq \Delta_{\min}$ for some Δ_{\min} . Subject to this assumption, we perform a “phase 1” optimization (described in section 6.2.2) to determine if a feasible spectrum exists. If one exists, we raise Δ_{\min} and repeat. If not, we lower Δ_{\min} and repeat.

In general, the most efficient way to find the maximal value of Δ_{\min} where a feasible spectrum exists is to perform a binary search. To determine Δ_ϵ within accuracy δ then requires $O(\log_2(1/\delta))$ runs of the simplex algorithm. This contrasts with p_T -maximization, which requires only two runs (phase 1 to determine an initial vertex, and a second run to perform the optimization). In addition, individual Δ_ϵ -maximization runs start taking more and more steps (hence longer and longer time) to find a feasible solution as Δ_{\min} approaches the boundary of the feasible region. On the other hand, p_T -maximization always works with a feasible spectrum and is free from such a slowdown. For these reasons, p_T -maximization is preferable.

³⁷These nearby pairs of dimensions are responsible for making the matrix \mathbf{A}_B nearly-degenerate. In technical language, its condition number decreases as the algorithm progresses. This is a well-known potential numerical instability of the semi-infinite programming problems, see [49], Figure 3. It is one of the reasons why we have to work with the multiple precision arithmetic—to avoid large rounding errors when inverting the poorly-conditioned matrix \mathbf{A}_B . Another reason for using multiple precision is not related to the “spectrum doubling”—it comes from the fact that for large N there is a huge disparity in size between low and high-order derivative components in the vectors $\mathbf{F}_{\Delta,\ell}^{\Delta_\sigma}$ composing the matrix \mathbf{A} . Computing the minimum reduced cost, Eq. (6.35), is an arithmetic operation which mixes these components. Since the minimum reduced cost may become tiny at the final steps of the algorithm, we have to perform its computation at a sufficient number of digits to be able to determine it accurately.

6.3 Implementation and Running Time Details

A typical sequence of steps in a p_T -maximization bootstrap computation is as follows. One first specifies N and runs code which computes and stores conformal block derivative expansion (6.5) (a table of $d_{\ell,i}^{m,n}$ and of $p_{\ell,i}^{m,n}$ polynomial coefficients) using one of the two methods from section 6.1. This is not computationally intensive. For example, our Mathematica code takes 12 minutes on an 8 core iMac to produce and store the 90MB table for $N = 231$ at the 64-digit precision using the method from section 6.1.1, expanding up to $n = 120$ in Eq. (6.7).

One then picks a Δ_σ and runs separate code which loads the above table, produces from it a similar Δ_σ -specific table for the derivatives of the functions $\mathbf{F}_{\Delta,\ell}^{\Delta_\sigma}$, and carries out the simplex algorithm described in the previous section. This code outputs the maximal attained value of p_T for the given Δ_σ , and the corresponding solution to crossing. These computations *are* intensive, and would be too slow to perform in Mathematica. We have two independently developed versions of this code, one in Python³⁸ and another in C++.

We will give performance details for the Python code (C++ code performance is similar). Hotstarting the p_T -maximization from the MFT spectrum makes phase 1 practically negligible, with a running time of a few minutes at $N = 231$. Phase 2 takes about 55 hours at $N = 231$, about 24 hours at $N = 190$, and about 6 hours at $N = 153$. These times are for an average single-core process on our clusters (composed of standard ~ 3 GHz machines with a ~ 4 GB memory limit per core; our computations require < 1 GB), and for Δ_σ in the transition region near the 3d Ising point. This region near the 3d Ising point appears to be the most time-consuming, while away from it the p_T -maximization concludes even faster. Also in 2d, the p_T -maximization is about factor two faster for the same N .

In a given computation, we may spawn $O(100)$ single-core Python processes for a number of Δ_σ values in an interval of interest. We did not keep careful track of the total CPU time used to produce the results of this paper, but we estimate it as 2 - 3 single-core CPU years. In comparison, the less precise Monte Carlo computations of [25] took about 30 CPU years, and would need about 1000 CPU years to get the leading critical exponents at the accuracy that we achieved here.

7 Summary and Discussion

In this work we performed a precision study of the 3d conformal bootstrap in the vicinity of the solution corresponding to the 3d Ising model at criticality. We also performed a detailed comparison with the analogous results in 2d. The goals were to both perform a high precision determination of operator dimensions as well as to gain insights into why the 3d Ising solution is special in the space of unitary solutions to crossing symmetry. We believe that we have succeeded on both of these fronts.

³⁸We use Cython [50] to link the MPFR multiple precision arithmetic library [51], and to compile more computationally intensive parts of our Python code.

First, using the bootstrap combined with the conjecture that the 3d Ising CFT minimizes the central charge we have determined the leading operator dimensions $\{\Delta_\sigma, \Delta_\epsilon\}$ in the 3d Ising CFT at a level that is 2 - 3 times more precise than the previous record determinations, by Monte Carlo techniques. Moreover, we have made new precise determinations of the central charge c , the OPE coefficient $f_{\sigma\sigma\epsilon}^2$, the second \mathbb{Z}_2 -even scalar dimension $\Delta_{\epsilon'}$ and its OPE coefficient $f_{\sigma\sigma\epsilon'}^2$, as well as the second \mathbb{Z}_2 -even spin 2 dimension $\Delta_{T'}$ and OPE coefficient $f_{\sigma\sigma T'}^2$. Using our techniques we additionally obtain reasonable estimates for all low dimension ($\Delta \lesssim 13$) scalar and spin 2 operators in the \mathbb{Z}_2 -even spectrum. Moreover, we showed that the same procedure in 2d can accurately reproduce the known spectrum of the 2d Ising CFT to a high degree of precision. It is now a challenge for other methods to test our very precise new 3d predictions. The new approach to lattice radial quantization being developed in [33] seems to be promising in this respect. The entanglement renormalization methods [52] are also known to compute very well 2d Ising operator dimensions and OPE coefficients; would they be competitive for 3d?

Second, we observed that by following the boundary of the region allowed by crossing symmetry, there is a dramatic rearrangement of the operator spectrum as one crosses the “kink” associated with the Ising model. This rearrangement shows compelling numerical evidence that the solution corresponding to the 3d Ising model contains fewer operators than a generic solution to crossing symmetry, i.e. that certain operators decouple as one approaches the Ising point. Such behavior is similar to what happens in 2d, where the decoupling of operators can be understood in terms of null states of the Virasoro symmetry. While there is no clear symmetry interpretation of the decoupled states in 3d, we believe that this gives an important characterization of why the 3d Ising solution is special, going beyond the conjecture that the central charge is minimized.

Could it be that the critical 3d Ising model is, after all, exactly solvable? While at present we do not have a good idea of how to turn our numerical results into a systematic search for an exact solution, we hope that our results will stimulate theoretical thought in this direction.

An alternative point of view could be that the 3d Ising CFT is not exactly solvable, yet we found a very efficient method to solve it numerically, much better than any previously known technique. Both possibilities are interesting, and the future will show which one is true.

While this fundamental dilemma may have to await its resolution for some time, we can see several research directions which can be attacked immediately. It is clearly important to better understand the interpretation of the decoupling states. For this a precision study of the interpolating solution between 2d and 3d, going beyond our recent results [3] would be useful. It would also be good to perform a more careful analysis of higher spin operators, where one can compare with MC determinations as well as verify “Nachtmann’s theorem” [21] and recent analytic results for the asymptotic behavior at large spin [18, 19] (see section 3.7).

It is also important to extend our analysis beyond the single correlator $\langle\sigma\sigma\sigma\sigma\rangle$, where e.g. the correlator $\langle\sigma\sigma\epsilon\epsilon\rangle$ would allow us to study the \mathbb{Z}_2 -odd spectrum. Work in this

direction is ongoing [5]. More generally, including higher dimension external operators should boost sensitivity to higher dimension exchanged operators, since the exponential decoupling sets in later in this case (see [9]). It would also be useful to more systematically study constraints from different regions of cross ratio space, where e.g. it is expected that going to the lightcone limit should boost sensitivity to high spin operators with low twist (see note 32).

One could also incorporate additional symmetries into precision studies of the bootstrap. It would be interesting to perform a similar precision analysis of the spectrum of the $O(N)$ vector models for different values of N , extending the bootstrap results of [4]. In the XY universality class ($N = 2$) there exists an 8σ discrepancy between the lattice determination of the ν critical exponent and the direct experimental measurement at the λ -point of ^4He (see [53] for a review). Borel-resummed perturbation theory methods are not sufficiently precise to say who’s right and who’s wrong. It seems likely that the bootstrap should be able to do so.

It would also be interesting to consider other discrete symmetries (e.g., \mathbb{Z}_n , $n \geq 3$) via the bootstrap. One could also do similar spectrum studies in 4d or other higher dimensions, where one could also look for evidence of transitions in the solutions living on the boundary. In fact, similar bounds in 4d $\mathcal{N} = 1$ supersymmetric theories show evidence of a “kink” [12], and it would be very interesting to study the corresponding transition in detail. More generally, similar spectrum studies can be performed in systems with various amounts of supersymmetry across different dimensions.

It should be clear that this paper brings us one step closer to the dream of the conformal bootstrap—that one can solve theories using only basic inputs about symmetries, even in more than 2 dimensions. An extremely powerful approach is gradually emerging, which will give us precision access to strongly-coupled physics in systems where perturbation theory and $1/N$ expansions break down. Many of these systems have beautiful realizations in condensed matter and statistical physics. For some of us, this experimental connection is the best justification of this research program. For loftier thinkers, bootstrap results may provide a window into the world of strongly-coupled quantum gravity via the AdS/CFT correspondence. There is something for everyone in the bootstrap!

Acknowledgements

We are grateful to M. Hasenbusch, M. Henkel, D. Mouhanna and E. Vicari for the useful communications concerning their work. We are grateful to B. van Rees for the discussions of the interpolating solution. In addition, we thank N. Arkani-Hamed, C. Beem, A.L. Fitzpatrick, G. Fleming, H. Ooguri, H. Osborn, J. Kaplan, E. Katz, F. Kos, J. Maldacena, J. Penedones, L. Rastelli, N. Seiberg, and A. Zhiboedov for related discussions.

S.R. is grateful to the Samara Chernorechenskaya Scientific Center for their hospitality. S.R., D.S.D., and D.P. are grateful to KITP for their hospitality. We would also like to thank the organizers and participants of the Back to the Bootstrap 3 conference at CERN.

This research was supported in part by the National Science Foundation under Grant No. PHY11-25915. The work of S.E. was partially supported by the French ANR contract 05-BLAN-NT09-573739, the ERC Advanced Grant no. 226371 and the ITN programme PITN-GA-2009-237920. M.P. is supported by DOE grant DE-FG02-11ER41742. A.V is supported by DOE grant DE-AC02-05CH1123. The work of D.S.D. is supported by DOE grant number DE-SC0009988.

Computations for this paper were run on National Energy Research Scientific Computing Center supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH1123; on the CERN cluster; on the Aurora and Hyperion clusters supported by the School of Natural Sciences Computing Staff at the Institute for Advanced Study; on the Omega cluster supported by the facilities and staff of the Yale University Faculty of Arts and Sciences High Performance Computing Center; on the TED cluster of the Chemistry Department and High Energy Theory group at Brown University; and the Kelvin cluster at the C.E.A. Saclay funded by the European Research Council Advanced Investigator Grant ERCAdG228301. S.E. would like to thank D. Kosower for providing access to the Kelvin cluster.

A Comparison to Semidefinite Programming

An alternative optimization algorithm that has proved useful in bootstrap studies is semidefinite programming [12, 4]. Like our algorithm presented here, semidefinite programming avoids discretizing the possible operator dimensions Δ . The starting point is an approximation for linear functionals in terms of polynomials times positive functions,

$$\sum_{m,n} a_{mn} \partial_r^m \partial_\eta^n G_{\Delta,\ell}(r, \eta) \Big|_{r=r_*, \eta=1} \approx \frac{r_*^\Delta}{Q_\ell(\Delta)} \sum_{m,n} a_{mn} P_\ell^{m,n}(\Delta), \quad (\text{A.1})$$

where $P_\ell^{m,n}(\Delta)$ are polynomials and $Q_\ell(\Delta) = \prod_i (\Delta - \Delta_i)$. This approximation follows from Eq. (6.5). Positivity of the polynomial $\sum_{m,n} a_{mn} P_\ell^{m,n}(\Delta)$ can be encoded in terms of $r \times r$ positive semidefinite matrices, where

$$r = \max_{m,n} \left[\frac{1}{2} \deg P_\ell^{m,n}(\Delta) \right] = \left[K + \frac{1}{2}q \right]. \quad (\text{A.2})$$

Here, the number of derivatives satisfies $m+n \leq 2K$, as in section 6.2.3, and q is the number of poles Δ_i included in the rational approximation (A.1).

The most popular semidefinite program solvers used in the bootstrap are SDPA and its arbitrary precision version SDPA-GMP [54]. Their performance scales differently from the performance of our algorithm. In practice, the running time increases slowly as the number of crossing relations n_C is increased, but quickly as the degrees of the polynomial approximations r (K or q) are increased. (By contrast, our algorithm is relatively insensitive to the number of poles q .) This makes SDPA a good choice for studying theories with global symmetries, where we have different crossing relations for each tensor structure which can

appear in a four-point function [55]. (For example, theories with $O(N)$ symmetry have $n_C = 3$ for a four-point function of vectors. Theories with $SU(N)$ symmetry have $n_C = 6$ for a four-point function of anti-/fundamentals.) The overall dimension of the space of linear functionals is

$$d_\Lambda = n_C N = n_C \frac{K(K+1)}{2}. \quad (\text{A.3})$$

In [12], this was taken as high as $n_C = 6$ and $K = 11$, so $d_\Lambda = 396$. Each optimization for a problem of this size takes approximately 48 hours.

In this work, we have been interested in studying a single crossing relation $n_C = 1$ and exploring as many derivatives as possible. We have been able to reach $K = 21$, so that $d_\Lambda = 231$. Although this value of d_Λ is smaller than what was achieved with SDPA-GMP in the case of global symmetries, the number of derivatives $K = 21$ could be difficult to match with SDPA-GMP. As the number of derivatives K is increased, one must also increase the number of poles q to maintain numerical stability. Further, one should tune the SDPA-GMP parameters to ensure the solver uses good initial data and termination criteria. It will be important to explore whether this can be done in the future.

SDPA-GMP uses a primal-dual solution method, so in principle it could be used for precision spectrum studies similar to what we do here. This will be interesting to explore in future work.

References

- [1] S. El-Showk, M. F. Paulos, D. Poland, S. Rychkov, D. Simmons-Duffin, and A. Vichi, “Solving the 3D Ising Model with the Conformal Bootstrap,” *Phys.Rev.* **D86** (2012) 025022, [arXiv:1203.6064 \[hep-th\]](#).
- [2] S. El-Showk and M. F. Paulos, “Bootstrapping Conformal Field Theories with the Extremal Functional Method,” *Phys. Rev. Lett.* **111** (2012) 241601, [arXiv:1211.2810 \[hep-th\]](#).
- [3] S. El-Showk, M. F. Paulos, D. Poland, S. Rychkov, D. Simmons-Duffin, and A. Vichi, “Conformal Field Theories in Fractional Dimensions,” *Phys. Rev. Lett.*, to appear (2014), [arXiv:1309.5089 \[hep-th\]](#).
- [4] F. Kos, D. Poland, and D. Simmons-Duffin, “Bootstrapping the $O(N)$ Vector Models,” [arXiv:1307.6856 \[hep-th\]](#).
- [5] F. Kos, D. Poland, and D. Simmons-Duffin, “Bootstrapping Mixed Correlators,” to appear.
- [6] I. Heemskerk, J. Penedones, J. Polchinski, and J. Sully, “Holography from Conformal Field Theory,” *JHEP* **0910** (2009) 079, [arXiv:0907.0151 \[hep-th\]](#).
- [7] A. L. Fitzpatrick and J. Kaplan, “Unitarity and the Holographic S-Matrix,” *JHEP* **1210** (2012) 032, [arXiv:1112.4845 \[hep-th\]](#).

- [8] R. Rattazzi, V. S. Rychkov, E. Tonni, and A. Vichi, “Bounding scalar operator dimensions in 4D CFT,” *JHEP* **12** (2008) 031, [arXiv:0807.0004 \[hep-th\]](#).
- [9] D. Pappadopulo, S. Rychkov, J. Espin, and R. Rattazzi, “OPE Convergence in Conformal Field Theory,” *Phys.Rev.* **D86** (2012) 105043, [arXiv:1208.6449 \[hep-th\]](#).
- [10] D. Poland and D. Simmons-Duffin, “Bounds on 4D Conformal and Superconformal Field Theories,” *JHEP* **1105** (2011) 017, [arXiv:1009.2087 \[hep-th\]](#).
- [11] R. Rattazzi, S. Rychkov, and A. Vichi, “Central Charge Bounds in 4D Conformal Field Theory,” *Phys. Rev.* **D83** (2011) 046011, [arXiv:1009.2725 \[hep-th\]](#).
- [12] D. Poland, D. Simmons-Duffin, and A. Vichi, “Carving Out the Space of 4D CFTs,” *JHEP* **1205** (2012) 110, [arXiv:1109.5176 \[hep-th\]](#).
- [13] T. Nishioka and K. Yonekura, “On RG Flow of τ_{RR} for Supersymmetric Field Theories in Three-Dimensions,” *JHEP* **1305** (2013) 165, [arXiv:1303.1522 \[hep-th\]](#).
- [14] S. Rychkov, “Conformal Bootstrap in Three Dimensions?,” [arXiv:1111.2115 \[hep-th\]](#).
- [15] J. Maldacena and A. Zhiboedov, “Constraining Conformal Field Theories with A Higher Spin Symmetry,” *J.Phys.* **A46** (2013) 214011, [arXiv:1112.1016 \[hep-th\]](#).
- [16] M. Campostrini, A. Pelissetto, P. Rossi, and E. Vicari, “Improved high-temperature expansion and critical equation of state of three-dimensional ising-like systems,” *Phys. Rev. E* **60** (1999) 3526–3563, [arXiv:cond-mat/9905078 \[cond-mat\]](#).
- [17] K. Wilson and J. B. Kogut, “The Renormalization group and the epsilon expansion,” *Phys.Rept.* **12** (1974) 75–200.
- [18] A. L. Fitzpatrick, J. Kaplan, D. Poland, and D. Simmons-Duffin, “The Analytic Bootstrap and AdS Superhorizon Locality,” *JHEP* **1312** (2013) 004, [arXiv:1212.3616 \[hep-th\]](#).
- [19] Z. Komargodski and A. Zhiboedov, “Convexity and Liberation at Large Spin,” *JHEP* **1311** (2013) 140, [arXiv:1212.4103 \[hep-th\]](#).
- [20] J. Callan, Curtis G. and D. J. Gross, “Bjorken scaling in quantum field theory,” *Phys.Rev.* **D8** (1973) 4383–4394.
- [21] O. Nachtmann, “Positivity constraints for anomalous dimensions,” *Nucl.Phys.* **B63** (1973) 237–247.
- [22] R. Guida and J. Zinn-Justin, “Critical exponents of the N vector model,” *J.Phys.* **A31** (1998) 8103–8121, [arXiv:cond-mat/9803240 \[cond-mat\]](#).
- [23] M. Campostrini, A. Pelissetto, P. Rossi, and E. Vicari, “25th-order high-temperature expansion results for three-dimensional ising-like systems on the simple-cubic lattice,” *Phys. Rev. E* **65** (2002) 066127, [arXiv:cond-mat/0201180 \[cond-mat\]](#).

- [24] Y. Deng and H. W. J. Blöte, “Simultaneous analysis of several models in the three-dimensional ising universality class,” *Phys. Rev. E* **68** (2003) 036125.
- [25] M. Hasenbusch, “Finite size scaling study of lattice models in the three-dimensional ising universality class,” *Phys. Rev. B* **82** (2010) 174433, [arXiv:1004.4486 \[cond-mat\]](#).
- [26] L. Canet, B. Delamotte, D. Mouhanna, and J. Vidal, “Nonperturbative renormalization group approach to the Ising model: A Derivative expansion at order ∂^4 ,” *Phys.Rev.* **B68** (2003) 064421, [arXiv:hep-th/0302227 \[hep-th\]](#).
- [27] D. F. Litim and D. Zappala, “Ising exponents from the functional renormalisation group,” *Phys.Rev.* **D83** (2011) 085009, [arXiv:1009.1948 \[hep-th\]](#).
- [28] K. E. Newman and E. K. Riedel, “Critical exponents by the scaling-field method: The isotropic $O(N)$ -vector model in three dimensions,” *Phys. Rev. B* **30** (1984) 6615–6638.
- [29] J. Comellas and A. Travesset, “O (N) models within the local potential approximation,” *Nucl.Phys.* **B498** (1997) 539–564, [arXiv:hep-th/9701028 \[hep-th\]](#).
- [30] D. F. Litim, “Critical exponents from optimized renormalization group flows,” *Nucl.Phys.* **B631** (2002) 128–158, [arXiv:hep-th/0203006 \[hep-th\]](#).
- [31] D. F. Litim and L. Vergara, “Subleading critical exponents from the renormalization group,” *Phys.Lett.* **B581** (2004) 263–269, [arXiv:hep-th/0310101 \[hep-th\]](#).
- [32] M. Henkel, “Finite Size Scaling and Universality in the (2+1)-dimensions Ising Model,” *J.Phys.* **A20** (1987) 3969.
- [33] R. Brower, G. Fleming, and H. Neuberger, “Lattice Radial Quantization: 3D Ising,” *Phys.Lett.* **B721** (2013) 299–305, [arXiv:1212.6190 \[hep-lat\]](#).
- [34] M. Weigel and W. Janke, “Universal amplitude ratios in finite-size scaling: three-dimensional ising model,” *Nuclear Physics B (Proceedings Supplements)* **83** (2000) 721, [arXiv:cond-mat/0009032](#).
- [35] S. J. Hathrell, “Trace Anomalies and $\lambda\phi^4$ Theory in Curved Space,” *Ann. Phys.* **139** (1982) 136.
- [36] I. Jack and H. Osborn, “Background Field Calculations in Curved Space-Time. 1. General Formalism and Application to Scalar Fields,” *Nucl. Phys.* **B234** (1984) 331.
- [37] A. Cappelli, D. Friedan, and J. I. Latorre, “C theorem and spectral representation,” *Nucl.Phys.* **B352** (1991) 616–670.
- [38] A. Petkou, “Conserved Currents, Consistency Relations, and Operator Product Expansions in the Conformally Invariant $O(N)$ Vector Model,” *Annals Phys.* **249** (1996) 180–221, [arXiv:hep-th/9410093](#).

- [39] J. Le Guillou and J. Zinn-Justin, “Accurate critical exponents for Ising like systems in noninteger dimensions,” *J. Physique* **48** (1987) 19–24.
- [40] A. Codello, “Scaling Solutions in Continuous Dimension,” *J.Phys.* **A45** (2012) 465006, [arXiv:1204.3877 \[hep-th\]](#).
- [41] V. S. Rychkov and A. Vichi, “Universal Constraints on Conformal Operator Dimensions,” *Phys. Rev.* **D80** (2009) 045006, [arXiv:0905.2211 \[hep-th\]](#).
- [42] A. Vichi, *A New Method to Explore Conformal Field Theories in Any Dimension*. Ph.D. Thesis, EPFL, 2011, 164 pp. <http://library.epfl.ch/theses/?nr=5116>.
- [43] P. Liendo, L. Rastelli, and B. C. van Rees, “The Bootstrap Program for Boundary CFT_d,” *JHEP* **1307** (2013) 113, [arXiv:1210.4258 \[hep-th\]](#).
- [44] D. Friedan, Z.-a. Qiu, and S. H. Shenker, “Conformal Invariance, Unitarity and Two-Dimensional Critical Exponents,” *Phys.Rev.Lett.* **52** (1984) 1575–1578.
- [45] D. Friedan, S. H. Shenker, and Z.-a. Qiu, “Details of the Nonunitarity Proof for Highest Weight Representations of the Virasoro Algebra,” *Commun.Math.Phys.* **107** (1986) 535.
- [46] M. Hogervorst and S. Rychkov, “Radial Coordinates for Conformal Blocks,” *Phys.Rev.* **D87** (2013) 106004, [arXiv:1303.1111 \[hep-th\]](#).
- [47] M. Hogervorst, H. Osborn, and S. Rychkov, “Diagonal Limit for Conformal Blocks in d Dimensions,” *JHEP* **1308** (2013) 014, [arXiv:1305.1321 \[hep-th\]](#).
- [48] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*. Cambridge, UK: Univ. Pr., 2007. 1235 p.
- [49] R. Reemtsen and S. Görner, “Numerical methods for semi-infinite programming: A survey,” in *Semi-Infinite Programming*, R. Reemtsen and J.-J. Rückmann, eds., pp. 195–275. Springer, 1998.
- [50] *Cython*. <http://cython.org>.
- [51] *MPFR*. <http://www.mpfr.org>.
- [52] G. Vidal, “Entanglement Renormalization: an introduction,” [arXiv:0912.1651 \[cond-mat.str-el\]](#).
- [53] E. Vicari, “Critical phenomena and renormalization-group flow of multi-parameter Φ^4 field theories,” *PoS LAT2007* (2007) 023, [arXiv:0709.1014 \[hep-lat\]](#).
- [54] *SDPA*. <http://sdpa.sourceforge.net/>.
- [55] R. Rattazzi, S. Rychkov, and A. Vichi, “Bounds in 4D Conformal Field Theories with Global Symmetry,” *J. Phys.* **A44** (2011) 035402, [arXiv:1009.5985 \[hep-th\]](#).