

Causal bootstrap: Bounding smeared spectral functions from nonperturbative Euclidean data

Ryan Abbott^{1,*}, Sarah Fields¹, William I. Jay², Patrick Oare³, and Matteo Saccardi²

¹*Physics Department, Columbia University, New York, New York 10027, USA*

²*Department of Physics, Colorado State University, Fort Collins, Colorado 80523, USA*

³*Physics Department, Brookhaven National Laboratory, Upton, New York 11973, USA*



(Received 10 June 2026; accepted 27 August 2026; published 16 October 2026)

This work introduces the causal bootstrap, a framework for bounding smeared spectral observables from finite nonperturbative Euclidean data. The method optimizes over the convex set of positive spectral densities compatible with the data to compute rigorous upper and lower bounds on the target observable. Statistical uncertainties, including correlations, are incorporated through compatibility regions using the covariance matrix. The bounds are equivalent, via Lagrange duality, to certified bounds on the target smearing kernel. For polynomial, rational, and piecewise rational kernels, the resulting dual problems can be reduced to finite-dimensional semidefinite programs using techniques familiar, e.g., in the numerical conformal bootstrap. The present formulation clarifies the relation to moment problems, Nevanlinna-Pick interpolation, and linear kernel-reconstruction methods. Numerical examples demonstrate the method.

DOI: [10.1103/m9kj-y6t3](https://doi.org/10.1103/m9kj-y6t3)

I. INTRODUCTION

Extracting spectral information from Euclidean-time correlation functions is a central problem in lattice field theory. In the simplest zero-temperature case, the full problem corresponds mathematically to inverting the Laplace transform

$$C(t) = \int_0^\infty dE \rho(E) e^{-Et}, \quad (1)$$

where $C(t)$ is a Euclidean-time connected correlation function and $\rho(E)$ is a spectral density. Euclidean correlators can be computed nonperturbatively from first-principles lattice field theory using Monte Carlo methods. These methods give $C(t)$ at a finite, discrete set of points with finite statistical uncertainty, rendering the reconstruction problem ill-posed.

In recent years significant progress has been made by avoiding direct reconstruction of the spectral function and instead targeting *smeared* spectral quantities [1–7]. That is, rather than attempting to determine $\rho(E)$ pointwise, one chooses a kernel $K(E)$ and attempts to compute an integrated quantity of the form

$$\mathcal{K}[\rho] = \int_0^\infty dE K(E) \rho(E). \quad (2)$$

Many phenomenologically relevant hadronic observables can be recast as smeared spectral reconstruction problems for different choices of K , including contributions to the anomalous magnetic moment of the muon, inclusive cross sections, transport coefficients, and inclusive decay rates [8–14]. Linear methods have been proposed to estimate such quantities from Euclidean data, including Backus-Gilbert [15,16] and the closely related approach of Hansen, Lupo, and Tantaló (HLT) [5], as well as methods based on Chebyshev polynomials [11,12,17], Gaussian processes [18–22], and the Mellin transform [23]. Nonlinear methods based on analyticity, moment problems, or convex optimization have also been proposed to obtain rigorous constraints on the allowed values of smeared spectral functions [24–29].

The central question addressed in this work is the following. Given a finite set of Euclidean correlator data and imposing explicitly the positivity of the spectral density, “What are the sharp upper and lower bounds on $\mathcal{K}[\rho]$?” This framing of the question turns spectral reconstruction into an optimization problem. For noisy Monte Carlo data, compatibility means lying within an uncertainty ellipsoid (induced by the covariance matrix) surrounding the measured correlator, leading to conservative confidence intervals for the target observable. In the limit of exact data, the ellipsoid shrinks and compatibility turns into an interpolation of the known correlator values.

*Contact author: rwa2110@columbia.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI. Funded by SCOAP³.

The present work develops this point of view as a bootstrap method for spectral functions. The term “bootstrap” is used here in the same broad sense as in the numerical conformal bootstrap [30]: a finite set of non-perturbative input data is combined with positivity constraints to produce rigorous bounds on quantities of interest. This use of “bootstrap” is distinct from the statistical bootstrap, wherein independent data are resampled in order to assess statistical uncertainties. The proposed method also differs from recent work on the quantum mechanical bootstrap (e.g., Refs. [31–34]) which incorporates positivity but do not leverage additional non-perturbative input data.

The first novel contribution of this work is the unification of several spectral reconstruction methods as a bootstrap procedure. Bounds coming from methods such as Nevanlinna-Pick (NP) interpolation [24] and moment problems [26,28] can be cast in the language of convex optimization problems discussed in Ref. [29], and such reformulations yield mathematically equivalent tight bounds on reconstructed smeared spectral quantities when respectively applied to equivalent input data, in a sense to be explained later (cf. Sec. V). These statements hold both for the scalar problems, as well as when generalized to Hermitian matrices of correlation functions. The matrix-valued formulation can improve bounds by leveraging additional operators, similar in spirit but wholly different in realization to the generalized eigenvalue problem (GEVP) approach to spectroscopy [35,36]. Methods from convex optimization are used to constrain the spectral function that produces the most extreme $\mathcal{K}[\rho]$ while remaining compatible with the observed data. The linear methods discussed in Refs. [5,11,12] are closely related to the dual formulation of the optimization problem pioneered in Ref. [29]. The primary distinction is that the latter introduces upper and lower bounds on the kernel, rather than directly targeting its controlled approximation. For positive spectral functions, this immediately allows one to infer rigorous bounds on the resulting smeared observable.

The second novel contribution of this work is to apply methods from the bootstrap literature in order to reduce the optimization problems discussed above to finite-dimensional semidefinite programs (SDPs) [37,38]. In particular, the basis functions used to construct bounds based on both Nevanlinna-Pick interpolation and moment problems admit a polynomial structure that allows any spectral reconstruction with a piecewise rational smearing kernel to be reduced to a mathematically equivalent finite-dimensional problem. This can then be solved efficiently using primal-dual interior point methods as in the numerical conformal bootstrap [37]. In contrast to the dual formulation discussed in Ref. [29], such formulations are entirely finite-dimensional, meaning that there are a finite number of variables and a finite number of constraints. Furthermore, bounds derived from primal-dual methods are

certified, in the sense that both the existence and tightness of the bounds can be explicitly verified via a short computation.¹

The remainder of this work is structured as follows. Sections II to IV develop the primal formulation, dual certificates, and finite SDP reduction. Section V establishes equivalences between certain families of bounds. Section VI presents numerical demonstrations of the proposed methods. Sections VII and VIII compare the method with related approaches and discuss practical assumptions for lattice applications.

Code for the numeric examples in Sec. VI is available at [41].

II. SPECTRAL RECONSTRUCTION PROBLEMS

This section defines the positive spectral measures, correlator functionals, and target observables used throughout the paper. These ingredients lead to primal optimization problems for exact and noisy data. Scalar versions of certain optimization problems have already appeared in the literature; Ref. [29] has discussed the scalar version of Eqs. (18) and (21). To the best of the authors’ knowledge, Eq. (23) is novel to the present work.

A. Spectral densities and smeared observables

Consider a set of operators $\{\mathcal{O}_0, \dots, \mathcal{O}_{r-1}\}$ in a zero-temperature quantum field theory. The associated matrix of two-point Euclidean correlation functions admits the spectral decomposition

$$C_{ab}(t) = \langle \mathcal{O}_a(t) \mathcal{O}_b(0)^\dagger \rangle = \int_0^\infty dE \rho_{ab}(E) e^{-Et}, \quad (3)$$

with the spectral density

$$\rho_{ab}(E) = \sum_{n=0}^\infty \langle 0 | \mathcal{O}_a | n \rangle \langle n | \mathcal{O}_b^\dagger | 0 \rangle \delta(E - E_n), \quad (4)$$

where $\{|n\rangle\}$ denotes an energy eigenbasis of the physical Hilbert space. From Eq. (4), it is clear that ρ is a positive semidefinite Hermitian matrix-valued spectral function; This well-known positivity can be understood as a consequence of reflection positivity, and is the guiding physical constraint used throughout this work. For lattice-regulated theories, this positivity follows from reflection positivity of the lattice action and the existence of a positive transfer matrix. When these conditions hold, the Euclidean

¹The terms *certified bound* and *certificate* are used here as is typical in the computer science literature to refer to an artifact or piece of data which can be used to construct a proof of a given statement, possibly with some additional computation. See Ref. [39], Sec. 5.5.1] for a discussion in convex optimization, and Ref. [40], Chap. 2] (or any computational complexity theory textbook) for a more broad discussion.

correlator admits a positive spectral representation of the form used in Eq. (4).

For the present applications, the target is a smeared spectral quantity, obtained by integrating the spectral density against a chosen kernel. For a Hermitian matrix-valued kernel $K(E)$, define

$$\mathcal{K}[\rho] = \int_0^\infty dE K(E) \bullet \rho(E), \quad (5)$$

where $X \bullet Y = \text{Tr}[X^\dagger Y]$ is the Hermitian trace pairing; for Hermitian X and Y , the trace pairing is $\text{Tr}[XY]$ and is real. More generally, the aim of this work is to determine the minimum and maximum possible values of Eq. (5) among all positive spectral densities compatible with the available Euclidean data.

B. General formulation

It will be useful to package several superficially distinct spectral reconstruction problems in a common notation. Throughout this formulation, matrix-valued quantities live in the space $\mathcal{H}^r$ of $r \times r$ Hermitian matrices. Let $I \subset \mathbb{R}$ be an interval, let r be a positive integer, and let $\mathcal{M}^r(I)$ denote the space of $r \times r$ Hermitian matrix-valued measures on I . Let $\mathcal{M}_+^r(I)$ denote the corresponding space of positive-semidefinite spectral functions. That is, $\rho \in \mathcal{M}_+^r(I)$ means that $\int_I dx f(x)^\dagger \rho(x) f(x) \geq 0$ for every compactly supported vector-valued test function $f: I \rightarrow \mathbb{C}^r$.

Given basis functions $b_t: I \rightarrow \mathbb{R}$, define the correlator data generated by ρ as

$$C_t[\rho] = \int_I dx b_t(x) \rho(x), \quad t = 0, \dots, N_t - 1. \quad (6)$$

Throughout the present work, C_t is assumed to be known on $t = 0, 1, \dots$, but the formulas straightforwardly generalize to any finite, discrete set. Unless indicated otherwise, all quantities are given in units of the lattice spacing. The left-hand side of Eq. (6) can be understood as a rank-3 tensor, $C_{tab} \equiv C_{ab}(t)$. For the zero-temperature spectral representation in Eq. (3), the interval is $I = [0, \infty)$ and the basis functions are $b_t(E) = e^{-Et}$. For a finite-temperature bosonic correlator with temporal extent N_t , the usual basis takes the form

$$b_t(E) = \frac{\cosh[E(t - N_t/2)]}{\sinh(EN_t/2)}. \quad (7)$$

The target observable is similarly written as

$$\mathcal{K}[\rho] = \int_I dx K(x) \bullet \rho(x), \quad (8)$$

where $K: I \rightarrow \mathcal{H}^r$ is the chosen smearing kernel.

The formulation in terms of basis functions also encompasses several analyticity-based approaches. For instance, moment problems arise from the zero-temperature spectral representation after the change of variables $\lambda = e^{-E}$. With $\lambda \in [0, 1]$, Eq. (3) becomes

$$C_{ab}(t) = \int_0^1 d\lambda \tilde{\rho}_{ab}(\lambda) \lambda^t, \quad (9)$$

where $\tilde{\rho}(\lambda(E)) = \lambda(E)^{-1} \rho(E)$. Thus the Euclidean correlator is the sequence of moments of a positive measure on $[0, 1]$, with basis functions

$$b_t^{\text{Moment}}(\lambda) = \lambda^t. \quad (10)$$

In the same variables, a smeared observable can be written as

$$\mathcal{K}[\rho] = \int_0^1 d\lambda K(-\log \lambda) \bullet \tilde{\rho}(\lambda), \quad (11)$$

up to the same change-of-variables convention used in the definition of $\tilde{\rho}$. This is the Hausdorff moment problem relevant for zero-temperature nonstaggered correlators [26].

Analyticity-based interpolation methods are naturally phrased in terms of the Stieltjes transform

$$G(z) = \int_I dx \frac{\rho(x)}{x - z}, \quad (12)$$

where z is a complex variable in the upper half plane and I is an interval containing the support of ρ . For a positive spectral density, G is a Pick function²: its imaginary part is positive semidefinite in the upper half plane. Conversely, the positive density can be recovered from the boundary value of the imaginary part through the Stieltjes-Perron inversion formula,

$$\rho(x) = \frac{1}{\pi} \lim_{\epsilon_0 \rightarrow 0} \text{Im } G(x + i\epsilon_0) \quad (13)$$

in the usual distributional sense. Together, Eqs. (12) and (13) provide a bijection between positive-semidefinite spectral functions and Pick functions; this is the key fact that allows analyticity-based approaches to be formulated in terms of a positive spectral function below. The imaginary part of $G(z)$ at a point $z = x_0 + i\epsilon_0$ is itself a smeared spectral function,

$$\frac{1}{\pi} \text{Im } G(x_0 + i\epsilon_0) = \int_I dx \rho(x) K_{x_0, \epsilon_0}^{\text{Cauchy}}(x), \quad (14)$$

²Pick functions are also equivalently called (Herglotz-)Nevanlinna functions in the literature.

where

$$K_{x_0, \epsilon_0}^{\text{Cauchy}}(x) = \frac{1}{\pi} \frac{\epsilon_0}{(x - x_0)^2 + \epsilon_0^2}. \quad (15)$$

If the input data are values of $G(z_n)$ at fixed points z_n in the upper half plane, their real and imaginary parts are again linear functionals of ρ . Indeed, taking real and imaginary parts of Eq. (12) gives the basis functions

$$b_n^{\text{NP,Re}}(x) = \frac{x - \text{Re} z_n}{(x - \text{Re} z_n)^2 + (\text{Im} z_n)^2}, \quad (16)$$

$$b_n^{\text{NP,Im}}(x) = \frac{\text{Im} z_n}{(x - \text{Re} z_n)^2 + (\text{Im} z_n)^2}, \quad (17)$$

where the basis functions are indexed by n instead of t to avoid confusion with Euclidean time. Thus, analyticity-based interpolation can be incorporated into the same primal problems, with a different choice of basis functions.

Once the basis functions and target kernel have been fixed, the goal is to compute the allowed range of $\mathcal{K}[\rho]$, constrained by spectral positivity and compatibility with the available data. The following two subsections make this statement precise, first for exact and then for noisy data.

C. Exact data

Suppose first that the correlator values $\hat{C}_t$ are known exactly for $t = 0, \dots, N_t - 1$. The lower bound on $\mathcal{K}[\rho]$ is obtained by minimizing the target observable over all positive densities that reproduce these data; see the following.

Problem II.1. (Primal spectral reconstruction problem for exact data).

$$\begin{aligned} & \underset{\rho \in \mathcal{M}_+^r(I)}{\text{minimize}} && \mathcal{K}[\rho] \\ & \text{subject to} && C_t[\rho] = \hat{C}_t \quad \forall t \in \{0, \dots, N_t - 1\}. \end{aligned} \quad (18)$$

The corresponding upper bound is obtained by maximizing $\mathcal{K}[\rho]$ subject to the same constraints, or equivalently by minimizing $-\mathcal{K}[\rho]$. Thus, a complete spectral bound requires solving two optimization problems, one for the lower end point and one for the upper end point.

Following the nomenclature of convex programming, Eq. (18) will be referred to as the primal problem. A choice of ρ is called *primal feasible* if it satisfies the positivity constraint and reproduces the correlator data. In the scalar case $r = 1$, Eq. (18) is an infinite-dimensional linear program. For $r > 1$, the Hermitian positivity condition makes it an infinite-dimensional semidefinite program. In the language of the moment-problem literature, the same problem is an instance of a generalized moment problem [42,43].

The infinite-dimensional nature of Eq. (18) presents a practical difficulty. Even neglecting finite-precision concerns and assuming that real numbers can be represented exactly, an arbitrary positive spectral function/distribution contains an infinite amount of information and hence cannot be represented directly on a computer. Truncating the space and introducing a discretization is possible, but would introduce an additional approximation. The purpose of Secs. III and IV is to show that, for important classes of basis functions and kernels, this infinite-dimensional problem can instead be transformed into a finite-dimensional semidefinite program without such discretization.

D. Inexact data

For practical lattice field theory calculations, Euclidean correlation functions are computed with finite statistical uncertainty. The interpolation constraint of Eq. (18) must therefore be relaxed to one of statistical consistency with the measured correlator.

Let $\hat{C}$ denote the rank-3 tensor of measured correlator data, $\hat{C}_{tab}$. The Hermitian, positive-definite covariance matrix Σ is a rank-6 tensor, $\Sigma_{tab,t'a'b'}$. The covariance matrix induces a norm defined by

$$\|C\|_\Sigma^2 = C^\dagger \Sigma^{-1} C = C_{tab}^* \Sigma_{tab,t'a'b'}^{-1} C_{t'a'b'}. \quad (19)$$

For numerical applications, it is convenient to use suitably vectorized versions of these quantities, in which case $\hat{C}_{tab}$ is understood to contain N_{corr} independent components. A correlator $C[\rho]$ is considered statistically consistent with $\hat{C}$ if

$$\|C[\rho] - \hat{C}\|_\Sigma \leq \sigma_0, \quad (20)$$

where $\sigma_0 > 0$ is a chosen cutoff. Equation (20) imposes a cutoff on the correlated χ^2 deviation between the candidate correlator and the measured central value. The cutoff may be chosen using an assumed Gaussian distribution, by a data-driven prescription, or by another procedure appropriate to the analysis [44,45].

The noisy-data version of the primal problem is as follows.

Problem II.2. (Primal spectral reconstruction problem for noisy data, second-order cone form).

$$\begin{aligned} & \underset{\rho \in \mathcal{M}_+^r(I)}{\text{minimize}} && \mathcal{K}[\rho] \\ & \text{subject to} && \|C[\rho] - \hat{C}\|_\Sigma \leq \sigma_0. \end{aligned} \quad (21)$$

As noted in Ref. [29], if the consistency region in Eq. (20) contains the true correlator with a given probability, then Eq. (21) gives bounds for $\mathcal{K}[\rho]$ at the same level of probability. They are conservative in the sense that the optimization propagates the entire allowed region of

correlators to the target observable. The resulting interval may therefore be wider than the uncertainty obtained from any particular estimator of $\mathcal{K}[\rho]$. Note that a solution to Eq. (21) need not always exist (e.g., for $\sigma_0 = 0$), since noisy data cannot be represented exactly by a positive spectral function. It may be advantageous to choose the cutoff σ_0 to be smaller than naive statistical interpretations would suggest. In particular, setting σ_0 equal to the number of degrees of freedom, naively estimated as the number of Euclidean data points, can yield statistical uncertainties that are more conservative than intended. However, consistent statistical interpretation of such a program introduces some subtleties, which is the subject of our ongoing future work.

Eq. (21) imposes a second-order cone constraint. For the finite-dimensional reduction below, it is useful to rewrite it as an equivalent semidefinite constraint. Using the Schur complement, Eq. (20) can be rewritten as

$$\Gamma[\rho] \equiv \begin{pmatrix} \Sigma & C[\rho] - \hat{C} \\ (C[\rho] - \hat{C})^\dagger & \sigma_0^2 \end{pmatrix} \succeq 0, \quad (22)$$

since Σ is positive definite. The same noisy-data problem can therefore be written as follows.

Problem II.3. (Primal spectral reconstruction problem for noisy data, SDP form).

$$\begin{aligned} & \underset{\rho \in \mathcal{M}_+^r(I)}{\text{minimize}} && \mathcal{K}[\rho] \\ & \text{subject to} && \Gamma[\rho] \succeq 0. \end{aligned} \quad (23)$$

The advantage of Eq. (23) is that the noisy-data constraint now has the same semidefinite character as the positivity condition on ρ . This observation will allow parallel treatment of exact and noisy data in the dual and finite-dimensional formulations below.

III. DUALITY

The primal problems in Eqs. (18) and (23) are formulated directly in terms of the spectral density. This section uses Lagrange duality to replace the optimization over positive spectral functions by a search for certified bounds on the target kernel. The dual problems have finitely many optimization variables, and their pointwise positivity constraints will be reduced to finite SDPs in Sec. IV. Scalar versions of some optimization problems have already appeared in the literature; the author of Ref. [29] has discussed the scalar version of Eq. (25). To the best of the authors' knowledge, Eqs. (36) and (41), as well as the application of the methods in Sec. III B to problems in lattice field theory are novel to the present work.

A. Lagrange duality and dual certificates for exact data

Lagrange duality provides a general procedure for converting a given primal optimization problem into an

equivalent dual problem. To apply Lagrange duality to the exact-data primal problem [Eq. (18)], introduce a positive multiplier $\mu \in \mathcal{M}_+^r(I)$ for the positivity constraint and matrices $g_t \in \mathcal{H}^r$ for the interpolation constraints. The Lagrange function is

$$\begin{aligned} L(\rho, \mu, g) &= \mathcal{K}[\rho] - \int_I dx \mu(x) \bullet \rho(x) - \sum_t g_t \bullet (C_t[\rho] - \hat{C}_t) \\ &= \int_I dx \rho(x) \bullet \left[K(x) - \sum_t g_t b_t(x) - \mu(x) \right] \\ &\quad + \sum_t g_t \bullet \hat{C}_t. \end{aligned}$$

The dual problem is obtained by minimizing L over unconstrained ρ , followed by maximizing over the Lagrange multipliers μ and g . Minimization over the unconstrained variable ρ is bounded from below only if

$$\mu(x) = K(x) - \sum_t g_t b_t(x) \geq 0, \quad \forall x \in I. \quad (24)$$

Eliminating μ delivers the dual problem shown below.

Problem III.1. (Dual spectral reconstruction, exact data).

$$\begin{aligned} & \underset{g_t \in \mathcal{H}^r}{\text{maximize}} && \sum_t g_t \bullet \hat{C}_t \\ & \text{subject to} && K(x) - \sum_t g_t b_t(x) \geq 0, \quad \forall x \in I. \end{aligned} \quad (25)$$

In this form, the dual problem can also be understood as an optimization problem over certified bounds on the primal problem. To see how, suppose that matrices $g_t \in \mathcal{H}^r$ can be chosen such that

$$M_g(x) \equiv K(x) - \sum_t g_t b_t(x) \geq 0, \quad \forall x \in I. \quad (26)$$

For reasons that will become clear, M_g will be referred to as the dual certificate. For any positive spectral density $\rho \in \mathcal{M}_+^r(I)$ compatible with the exact data, positivity implies

$$\begin{aligned} 0 &\leq \int_I dx M_g(x) \bullet \rho(x) \\ &= \mathcal{K}[\rho] - \sum_t g_t \bullet C_t[\rho] \\ &= \mathcal{K}[\rho] - \sum_t g_t \bullet \hat{C}_t. \end{aligned} \quad (27)$$

Thus, every choice of g_t satisfying Eq. (26) provides a rigorous lower bound

$$\mathcal{K}[\rho] \geq \sum_t g_t \cdot \hat{C}_t \quad (28)$$

for every feasible spectral density. In this sense, the dual variables define a one-sided reconstruction of the target kernel by the basis functions. The inequality in Eq. (26), rather than an approximate equality, is what turns this reconstruction into a certificate. The best lower bound obtainable in this way is the solution of the dual optimization problem.

A feasible point of Eq. (25) is called dual feasible. The optimal value of the dual problem is always a lower bound on the optimal value of the primal problem, a statement known as weak duality. Explicitly, if g^* is dual optimal and ρ_* is primal optimal, then

$$\mathcal{K}[\rho_*] \geq \sum_t g_t^* \cdot \hat{C}_t. \quad (29)$$

Under mild regularity assumptions, this inequality is saturated, and the primal and dual optima agree. This stronger statement, known as strong duality, is expected to hold in most physical applications by Slater's condition, which roughly requires the existence of a strictly feasible point [39,46]. An instance where Slater's condition would not apply arises in analyticity-based interpolation methods when the input data specify the feasible spectral density uniquely. However, such extremal problems occur with probability zero, and hence are not likely to cause issues in practice.

B. Extremal feasible measures

The dual problem does not directly return the physical spectral density. Instead, it identifies the most constraining one-sided approximation to the target kernel. Nevertheless, when strong duality holds, the optimal dual solution contains useful information about the spectral density that saturates the bound. This extremal density should not be interpreted as the true physical spectrum but rather as the feasible spectral density that produces the most extreme allowed value of the chosen observable. Similar methods to those described in this section have also seen use in the numerical conformal bootstrap, referred to as the *extremal functional method* [47,48].

Let $M_g(x)$ be the dual certificate defined in Eq. (26). At primal and dual optimum, saturation of Eq. (27) requires

$$\rho_*(x)M_g(x) = 0, \quad \forall x \in I \quad (30)$$

in the sense of positive matrix measures. In the convex-optimization literature, this condition is known as complementary slackness. In the scalar case, it implies that ρ_* can have support only where the dual certificate vanishes, i.e., $M_{g^*}(x) = 0$. If the zero set is a finite set $\{x_i\}_{i=1}^N$, the extremal feasible measure takes the form

$$\rho_*(x) = \sum_{i=1}^N A_i \delta(x - x_i), \quad A_i \geq 0. \quad (31)$$

The support points x_i are therefore determined by the zeros of the optimal dual certificate. The coefficients A_i are then constrained by the exact data,

$$C_t[\rho_*] = \sum_{i=1}^N A_i b_t(x_i) = \hat{C}_t, \quad (32)$$

and by saturation of the bound,

$$\mathcal{K}[\rho_*] = \sum_{i=1}^N A_i K(x_i) = \sum_t g_t^* \cdot \hat{C}_t. \quad (33)$$

Together these equations give finite-dimensional linear constraints on the non-negative coefficients A_i . These constraints can be solved using standard linear-algebra methods or non-negative least squares [49].

For matrix correlators, complementary slackness is less restrictive. A product of positive-semidefinite matrices can vanish even when neither factor vanishes. Thus, in the matrix case, Eq. (30) constrains the null directions of $M_{g^*}(x)$ rather than simply localizing the measure to zeros of a scalar function. Although the dual certificate is generically a full-rank object, $M_{g^*}(x)$ at the dual optimum is typically observed to be rank deficient in the numeric examples discussed below. The key point for the present work is that the extremal measure is a diagnostic of the bounding problem, not an independent extraction of the physical spectrum.

C. Noisy data

A similar dual formulation extends to the noisy-data primal problem in Eq. (23). The only new feature is that compatibility with the data is encoded by the semidefinite constraint $\Gamma[\rho] \geq 0$ rather than by exact interpolation conditions. The Lagrange multiplier for this constraint is a positive-semidefinite matrix $Z \in \mathcal{H}^{N_{\text{corr}}+1}$. Define, using Eq. (22), the quantity

$$\Gamma_0 \equiv \Gamma[\rho = 0] = \begin{pmatrix} \Sigma & -\hat{C} \\ -\hat{C}^\dagger & \sigma_0^2 \end{pmatrix}. \quad (34)$$

The entries of Z conjugate to the off-diagonal block of $\Gamma[\rho]$ play the role of the matrices g_t appearing in the exact-data dual. More explicitly, this embedding can be written schematically as

$$Z = \frac{1}{2} \begin{pmatrix} * & g \\ g^\dagger & * \end{pmatrix}, \quad (35)$$

where the ordering of the components of g matches the vectorization of the correlator data.

Dualizing Eq. (23) gives the following.

Problem III.2. (Dual spectral reconstruction, noisy data).

$$\begin{aligned} & \underset{Z \in \mathcal{H}^{N_{\text{corr}}+1}}{\text{maximize}} && -Z \bullet \Gamma_0 \\ & \text{subject to} && K(x) - \sum_t g_t b_t(x) \geq 0, \quad \forall x \in I, \\ & && Z \geq 0, \end{aligned} \quad (36)$$

with g_t embedded in Z as in Eq. (35). As in the exact-data case, the pointwise positivity condition is the kernel certificate. The new matrix constraint $Z \geq 0$ encodes the covariance ellipsoid that defined statistical consistency in the primal problem. Every feasible Z provides a rigorous lower bound

$$\mathcal{K}[\rho] \geq -Z \bullet \Gamma_0 \quad (37)$$

for every positive spectral density satisfying $\Gamma[\rho] \geq 0$.

The noisy-data dual has an equivalent form that makes contact with the exact-data certificate more directly. Starting from Eq. (21), one reformulates the constraint using a slack variable y :

$$\|C[\rho] - \hat{C}\|_{\Sigma} \leq \sigma_0 \Leftrightarrow \begin{cases} C[\rho] - \hat{C} - y = 0, \\ \|y\|_{\Sigma} \leq \sigma_0. \end{cases} \quad (38)$$

The second line on the right-hand side constrains the slack variable y to lie within an ellipsoid. The associated Lagrange function is

$$\begin{aligned} L(\rho, y; g, \mu) = & \mathcal{K}[\rho] - \int \mu(x) \bullet \rho(x) dx \\ & - g \bullet (C[\rho] - \hat{C} - y), \end{aligned} \quad (39)$$

where $\mu \in \mathcal{M}_+^r(I)$ enforces the positivity of ρ . Minimization over ρ is bounded from below only when the dual certificate from Eq. (26) satisfies $M_g(x) = \mu(x) \geq 0$ for all $x \in I$. Minimization of y over the ellipsoid gives

$$\underset{\|y\|_{\Sigma} \leq \sigma_0}{\text{minimize}} \quad g \bullet y = -\sigma_0 \sqrt{g^\dagger \Sigma g}. \quad (40)$$

Combining these results gives the dual problem as a second-order cone problem (SOCP) [50] as follows.

Problem III.3. (Dual spectral reconstruction, noisy data, SOCP form).

$$\begin{aligned} & \underset{g_t \in \mathcal{H}^r}{\text{maximize}} && \sum_t g_t \bullet \hat{C}_t - \sigma_0 \sqrt{g^\dagger \Sigma g} \\ & \text{subject to} && K(x) - \sum_t g_t b_t(x) \geq 0, \quad \forall x \in I. \end{aligned} \quad (41)$$

The first term is the exact-data dual objective evaluated on the measured central values. The second term penalizes directions in data space that are poorly determined by the covariance matrix. In the limit $\sigma_0 \rightarrow 0$, this penalty vanishes and Eq. (41) reduces to the exact-data dual problem. For the finite-dimensional SDP reduction in Sec. IV, it will be more convenient to use the SDP form Eq. (36). A similar dual formulation was reached in Ref. [29]; Eq. (41) is equivalent to Eq. (28) therein after optimization over the variable Ref. [29] refers to as μ .

The noisy-data problem also has complementary-slackness conditions, which are useful as diagnostics of extremal solutions. For any primal feasible ρ and dual feasible Z , the following equality holds:

$$\mathcal{K}[\rho] = -\Gamma_0 \bullet Z + \Gamma[\rho] \bullet Z + \int_I dx \rho(x) \bullet M_g(x), \quad (42)$$

where $M_g(x) = K(x) - \sum_t g_t b_t(x)$ is the dual certificate from Eq. (26). The last two summands on the right-hand side are non-negative for feasible primal and dual points. Thus equality at the optimum requires the following two conditions:

$$\rho_*(x) M_{g^*}(x) = 0, \quad \forall x \in I, \quad (43)$$

$$\Gamma[\rho_*] Z_* = 0. \quad (44)$$

The former is the same complementary slackness condition that appeared for exact data. The latter condition describes how the extremal correlator $C[\rho_*]$ sits on the boundary of the statistical consistency region. To see this, write

$$\Delta C[\rho] = C[\rho] - \hat{C}. \quad (45)$$

Taking the Schur complement of $\Gamma[\rho]$ shows that its lower-right Schur complement is

$$\sigma_0^2 - \Delta C[\rho]^\dagger \Sigma^{-1} \Delta C[\rho] \geq 0. \quad (46)$$

When the optimal dual matrix Z_* is nonzero and Σ is nonsingular, complementary slackness forces the inequality to be saturated in order to satisfy Eq. (44). That is, the extremal correlator lies on the boundary of the allowed ellipsoid,

$$\Delta C[\rho_*]^\dagger \Sigma^{-1} \Delta C[\rho_*] = \sigma_0^2. \quad (47)$$

This result makes manifest the fact that the primal optimizer will generically not reproduce the measured central values. Indeed, it identifies the positive spectral density whose correlator is statistically consistent with the data while extremizing the target observable.

IV. FINITE-DIMENSIONAL REDUCTION

The dual problems in Eqs. (25) and (36) have finitely many optimization variables but still contain infinitely many constraints: for every point x in the spectral domain, the dual certificate in Eq. (26) must be positive semidefinite. This section explains when those infinitely many positivity conditions can be replaced by a finite set of semidefinite constraints. The key point is that polynomial matrix positivity conditions in a single variable can be represented as a finite sum of squares [51]. Whenever the dual constraints can be written as polynomial matrix inequalities, the spectral reconstruction problem can be reduced to a finite-dimensional semidefinite program.

The formulation of spectral reconstruction problems Eqs. (49), (52), (53) and (56) as polynomial matrix problems that can be reduced to finite semidefinite programs is a central result of the present work.

A. Polynomial matrix programs

A polynomial matrix program (PMP) is a semidefinite program of the following form [37].

Problem IV.1. (General polynomial matrix program).

$$\begin{aligned} & \underset{y \in \mathbb{R}^n}{\text{maximize}} && b \cdot y \\ & \text{subject to} && M_j^0(x) + \sum_{k=1}^n y_k M_j^k(x) \geq 0, \\ & && \forall x \in I, \quad j = 1, \dots, J. \end{aligned} \quad (48)$$

Here $I \subseteq \mathbb{R}$ is an interval, the variables are $y_k \in \mathbb{R}$, and the matrices $M_j^k(x)$ are matrix-valued polynomials in x . Eq. (48) contains an affine matrix-valued constraint on the y_k . The same structure appears in spectral reconstruction whenever the target kernel and basis functions are polynomial, or can be made polynomial by multiplying by a positive denominator.

A simple example is the exact-data moment problem. On an interval I (e.g., $I = [0, 1]$) and with basis functions $b_t(\lambda) = \lambda^t$, the exact dual problem is shown below.

Problem IV.2. (Dual moment problem).

$$\begin{aligned} & \underset{g_t \in \mathcal{H}^r}{\text{maximize}} && \sum_t g_t \cdot \hat{C}_t \\ & \text{subject to} && K(\lambda) - \sum_t g_t \lambda^t \geq 0 \quad \forall \lambda \in I. \end{aligned} \quad (49)$$

If $K(\lambda)$ is a polynomial matrix, then Eq. (49) is already a PMP after expanding the Hermitian matrices g_t in a basis. Let $\{T_\alpha\}$ be a basis for $\mathcal{H}^r$, with $\alpha = 1, \dots, r^2$, and write

$$g_t = \sum_\alpha y_{t\alpha} T_\alpha. \quad (50)$$

Substituting Eq. (50) into Eq. (49) gives a polynomial matrix inequality that is affine in the variables $y_{t\alpha}$, and hence has the form of Eq. (48).

Many kernels of physical interest are more naturally represented as rational functions than as polynomials. This structure is still compatible with the PMP reduction, provided the denominator is positive on the interval where the constraint is imposed. For example, suppose that in Eq. (49) the kernel takes the form

$$K(\lambda) = \frac{n(\lambda)}{d(\lambda)} e_0 e_0^\top, \quad (51)$$

where $n, d \in \mathbb{R}[\lambda]$, $d(\lambda) > 0$ for all $\lambda \in I$, and $e_0 = (1, 0, \dots, 0)$. Such kernels arise naturally when reconstructing a smeared spectral function for the first interpolating operator $\mathcal{O}_0$, with the remaining operators serving to capture additional spectral information. The contributions of other operators can be isolated similarly. Multiplying the dual constraint by $d(\lambda)$ preserves the order $\geq$ and gives the equivalent PMP.

Problem IV.3. (Dual moment problem, rational kernel).

$$\begin{aligned} & \underset{g_t \in \mathcal{H}^r}{\text{maximize}} && \sum_t g_t \cdot \hat{C}_t \\ & \text{subject to} && n(\lambda) e_0 e_0^\top - \sum_t g_t \lambda^t d(\lambda) \geq 0 \quad \forall \lambda \in I. \end{aligned} \quad (52)$$

Exact knowledge of $G(z)$ as in Eq. (14) instead allows a similar treatment for a different class of kernels, namely those which are rational in E , rather than in $\lambda = e^{-E}$. Using the basis functions in Eqs. (16) and (17), the exact-data dual has the following form.

Problem IV.4. (Nevanlinna-Pick-based interpolation, rational form).

$$\begin{aligned} & \underset{g_n^{\text{Re}}, g_n^{\text{Im}}}{\text{maximize}} && \sum_n (g_n^{\text{Re}} \cdot \text{Re} \hat{G}_n + g_n^{\text{Im}} \cdot \text{Im} \hat{G}_n) \\ & \text{subject to} && K(E) - \sum_n \frac{(E - \text{Re} z_n) g_n^{\text{Re}} + (\text{Im} z_n) g_n^{\text{Im}}}{(E - \text{Re} z_n)^2 + (\text{Im} z_n)^2} \geq 0, \\ & && \forall E \in I. \end{aligned} \quad (53)$$

When $\text{Im} z_n > 0$, each denominator is positive on the real axis. Multiplying by the product of the denominators therefore converts the constraint into a polynomial matrix inequality, provided the target kernel is polynomial or rational with a compatible positive denominator. The price of this exact conversion is that the polynomial degree can become large as the number of interpolation points grows.

For kernels that are not polynomial or rational in the variable used by the basis functions, one may approximate the target kernel by a polynomial or rational function. In fact, the Stone-Weierstrass theorem, which says that a continuous function can be approximated uniformly by a

polynomial, suggests that such approximations will often converge rapidly. Regardless, such an approximation changes the target observable and should be treated as a separate source of systematic uncertainty. The finite-dimensional reduction is exact for the approximating kernel, but not automatically for the original kernel. In applications, the resulting error must therefore be separately bounded or otherwise included in the final uncertainty budget.

In contradistinction to existing linear methods such as those of Refs. [5,12], the degree of approximation in the current proposal is *not* directly tied to the number of input data used. However, for a fixed-order polynomial approximation to a given kernel, the *strength* of the bounds coming from the dual certificate [Eq. (26)] is generically expected to vary strongly depending on the number and quality of the input data.

B. Finite SDP reduction

Once the dual problem has been written as a PMP, the remaining task is to replace the constraint in Eq. (48), $M(x) \geq 0$ for all $x \in I$, by a finite number of semidefinite constraints. For polynomial matrices of a single variable, this reduction is possible using sum-of-squares representations. The precise form depends on whether I is the full real line, a half line, or a bounded interval. The details are collected in Appendix A.

Conceptually, the reduction works as follows. Choose a polynomial basis $\{q_i(x)\}$ and form the matrix of products $Q_{ij}(x) = q_i(x)q_j(x)$. Positivity of a polynomial matrix $M(x)$ on an interval can be represented by positive-semidefinite coefficient matrices Y_a such that

$$M(x) = \text{a finite linear expression in } Y_a \text{ and } Q(x). \quad (54)$$

Equating this expression to the polynomial matrix appearing in the PMP gives finitely many linear equations in the variables y_k and the positive matrices Y_a . The result is an ordinary finite-dimensional SDP.

For the exact-data problem, the final SDP has the following structure:

$$\begin{aligned} & \underset{y}{\text{maximize}} && b \cdot y \\ & \text{subject to} && \text{Tr}[A_p Y] + (By)_p = c_p, \quad p = 1, \dots, P, \\ & && Y \geq 0. \end{aligned} \quad (55)$$

The matrices A_p , the linear map B , and the constants c_p are determined directly by the polynomial identities used to enforce positivity. Values for A_p , B , and c_p can be found either by matching polynomial coefficients between Eqs. (54) and (55) or by evaluating the polynomial identity at sufficiently many sample points. Additional information on obtaining A_p , B , and c_p is given in Sections 2.2 and 2.5.2 of Ref. [37]. Different interval choices and block

decompositions change these matrices, but not the overall structure of the problem.

C. Extensions: Noisy data and piecewise kernels

The noisy-data dual problem in Eq. (36) fits into the same framework. The pointwise positivity constraint on the dual residual $K(x) - \sum_i g_i b_i(x)$ is reduced exactly as in the exact-data problem. The only additional ingredient is the semidefinite matrix variable Z that encodes the covariance ellipsoid. After the PMP constraints are reduced to finite form, the resulting SDP contains both the sum-of-squares matrix variable Y and the noisy-data matrix variable Z . The result is as follows.

Problem IV.5. (Dual spectral reconstruction, finite SDP).

$$\begin{aligned} & \underset{Z}{\text{maximize}} && -Z \cdot \Gamma_0 \\ & \text{subject to} && \text{Tr}[A_p Y] + (By)_p = c_p, \quad p = 1, \dots, P, \\ & && \text{Tr}[\hat{A}_{t\alpha} Z] - y_{t\alpha} = 0, \\ & && t = 0, \dots, N_t - 1, \quad \alpha = 1, \dots, r^2, \\ & && Y \geq 0, \quad Z \geq 0. \end{aligned} \quad (56)$$

The $y_{t\alpha}$ are the coefficients of g_t in a real basis of $\mathcal{H}^r$, while the $\hat{A}_{t\alpha}$ isolate the entries of Z as in Eq. (35). More explicitly, the matrices $\hat{A}_{t\alpha}$ can be defined by the relation

$$\sum_{\alpha} \text{Tr} \left[\hat{A}_{t\alpha} \begin{pmatrix} * & g \\ g^\dagger & * \end{pmatrix} \right] T_{ab}^\alpha = 2g_{t,ab}. \quad (57)$$

The matrices Y and Z may also be combined into a larger block-diagonal semidefinite variable $Y' = \text{diag}(Y, Z)$. With a corresponding block-diagonal objective matrix $D = \text{diag}(0, -\Gamma_0)$, the noisy-data problem fits the more general SDP form

$$\begin{aligned} & \underset{y, Y'}{\text{maximize}} && D \cdot Y' \\ & \text{subject to} && \text{Tr}[A'_p Y'] + (B'y)_p = c'_p, \quad p = 1, \dots, P', \\ & && Y' \geq 0, \end{aligned} \quad (58)$$

where A'_p, B', c'_p are appropriate combinations of the unprimed variables above. In implementation, the block structure of Y' should be preserved rather than treated as a dense matrix. This is the same kind of block structure exploited in conformal-bootstrap SDP solvers. The block structure may be important computationally but does not change the bound.

The same logic also handles kernels that are polynomial or rational only piecewise, which is particularly useful for observables with kinematic thresholds. For example, inclusive decays often involve kernels of the form

$$K_{\text{thresh}}(E) = \theta(M - E)\tilde{K}(E), \quad (59)$$

where $\theta(\cdot)$ is a Heaviside step function, M is a threshold energy, and $\tilde{K}(E)$ is polynomial or rational on the relevant interval. Instead of approximating the step function, the dual positivity condition can be split across the intervals on which the kernel has different analytic forms. For $I = [0, \infty)$, the dual constraint

$$K_{\text{thresh}}(E) - \sum_t g_t b_t(E) \geq 0 \quad \forall E \in [0, \infty) \quad (60)$$

is equivalent to the pair of constraints

$$\tilde{K}(E) - \sum_t g_t b_t(E) \geq 0, \quad E \in [0, M], \quad (61a)$$

$$-\sum_t g_t b_t(E) \geq 0, \quad E \in [M, \infty). \quad (61b)$$

Each interval produces its own sum-of-squares block in the finite SDP. Thus, piecewise polynomial or rational kernels can be handled exactly at the level of the dual constraints, provided each piece is polynomial or rational in the variable used for the PMP reduction.

After the PMP reduction, the remaining numerical task is to solve a finite-dimensional SDP. This work follows the primal-dual interior point strategy used in conformal-bootstrap calculations and in SDPB [37,52]. Such methods solve primal and dual finite SDPs simultaneously. The corresponding duality gap is computable and provides a certificate that the finite SDP has been solved to the desired precision.

V. EQUIVALENCE OF BOUNDS BETWEEN DIFFERENT FORMULATIONS

Once seemingly different problems have been expressed in the common language of convex optimization, it becomes clear that their formulations are in fact equivalent. This equivalence is a central result of the present work.

For instance, let $\{\hat{C}_t\}$ be a set of moments specifying a nondegenerate truncated Hamburger moment problem with $I = \mathbb{R}$ and $b_t(\lambda) = \lambda^t$, and consider the problem of determining the *Wertevorrat* defined by

$$\Delta_n(z) = \{G(z) | C_t[\rho] = \hat{C}_t\}, \quad (62)$$

$G(z)$ is the Stieltjes transform of the spectral function ρ ; cf. Eq. (12). Using the methods derived from the moment problem and Nevanlinna-Pick literature discussed in Refs. [24,26,28], it is possible to determine $\Delta_n(z)$ to be a circle with analytically calculable center and radius as

functions of the known moments $\hat{C}_t$ only. The derivation of the *Wertevorrat* is deeply steeped in the analytic properties of $G(z)$, and it is therefore remarkable that the same information can be accessed via an SDP formulation that references only the positivity of the spectral function with a given support.

To see this fact concretely, consider for simplicity a scalar $G(z)$ and define the following family of linear functionals on ρ parametrized by $\theta \in [0, 2\pi)$:

$$\mathcal{K}_\theta[\rho] = \text{Re}[e^{i\theta}G(z)] = \int_I d\lambda \rho(\lambda) K_\theta(\lambda), \quad (63)$$

$$K_\theta(\lambda) = \text{Re}\left(\frac{e^{i\theta}}{z - \lambda}\right). \quad (64)$$

The parameter θ allows one to sweep over the real and imaginary parts of $G(z)$ in order to map out the complete allowed region. Examples of equivalent bounds on $G(z)$ can now be stated precisely.

Lemma 1. (Equivalence of moment bounds). Consider a nonextremal truncated scalar Hamburger moment problem ($\lambda \in \mathbb{R}$) with moments $\hat{C}_t$, where $t = 0, \dots, N = 2n$. For a given fixed z in the upper half plane, $G(z)$ is constrained to lie within a convex set

$$\Delta_N(z) = \{G(z) | \rho \geq 0 \text{ has moments } \hat{C}_0, \dots, \hat{C}_N\}. \quad (65)$$

The following two descriptions for $\Delta_N(z)$ are equivalent:

- (1) $\Delta_N(z)$ is a disk with known analytic functions for the center $c_N(z)$ and radius $r_N(z)$ in terms of the input moments $\hat{C}_t$, given explicitly in Refs. [26,53].
- (2) $\Delta_N(z)$ is the intersection of the bounds obtained from solutions of the semidefinite programs for $\theta \in [0, 2\pi)$

$$\begin{aligned} &\text{minimize}_{\rho \in \mathcal{M}_+(\mathbb{R})} \mathcal{K}_\theta[\rho] \\ &\text{subject to } C_t[\rho] = \hat{C}_t, \quad t = 0, \dots, N. \end{aligned} \quad (66)$$

A proof of Lemma 1 is sketched in Appendix C. For matrix-valued problems, the set $\Delta_N(z)$ is generically a Weyl matrix ball. Similar results hold for moment problems over restricted intervals provided the interval is imposed in both setups.

The equivalences also directly extend to the dual formulation. For instance, Nevanlinna-Pick interpolation can be reframed as a dual optimization problem as follows.

Lemma 2. (Equivalence of interpolation bounds). Consider a nonextremal complex interpolation problem in which $G(z)$ is specified at a finite set of interpolation nodes $\{z_n\}$ in the upper half plane, $G(z_n) = \hat{G}_n$ for

$n = 0, \dots, N$. For a given $z \notin \{z_n\}$ in the upper half plane, $G(z)$ is constrained to lie within a convex set:

$$\Delta_N(z) = \{G(z) | \rho \in \mathcal{M}_+^1([0, \infty)), G(z_n) = \hat{G}_n\}. \quad (67)$$

The following descriptions for $\Delta_N(z)$ are equivalent:

- (1) $\Delta_N(z)$ is the intersection of the two disks obtained by solving the Nevanlinna-Pick interpolation problem on the nodes $\{z_n\}$ for the functions $G(z)$ and $zG(z)$ [54]. Explicit analytic formulas for $\Delta_N(z)$ are given in Ref. [24] and the references therein.
- (2) $\Delta_N(z)$ is the interior of the intersection of bounds obtained from solutions of the semidefinite programs for $\theta \in [0, 2\pi)$

$$\begin{aligned} & \underset{g_n^{\text{Re}}, g_n^{\text{Im}}}{\text{maximize}} \quad \sum_n (g_n^{\text{Re}} \text{Re} \hat{G}_n + g_n^{\text{Im}} \text{Im} \hat{G}_n) \\ & \text{subject to} \quad K_\theta(E) - \sum_n \frac{(E - \text{Re} z_n) g_n^{\text{Re}} + (\text{Im} z_n) g_n^{\text{Im}}}{(E - \text{Re} z_n)^2 + (\text{Im} z_n)^2} \\ & \quad \geq 0, \quad \forall E \in [0, \infty). \end{aligned} \quad (68)$$

The proof for Lemma 2 is very similar to that of Lemma 1 and thus omitted. Similar equivalences are expected to hold for arbitrary intervals $I \subseteq \mathbb{R}$ provided the restriction to I is included consistently on both sides. For the bounds coming from moment problems, the restriction is generically expected to be phrased in terms of intersections of Weyl matrix balls as in Ref. [55].

Note that the SDP and analyticity-based formulations here provide complementary approaches to the same problem: the analyticity-based approaches do not require an optimization step, and directly give the full structure of the Wertevorrat, whereas the SDP formulation readily generalizes to any polynomial or rational kernel and to noisy data. Which approach is more useful for a given calculation should be determined on a case-by-case basis.

VI. NUMERICAL RESULTS

This section presents numerical demonstrations of the bounds developed above. The examples below use a custom Python implementation with arbitrary-precision arithmetic from MPMATH [56]. Implementation details, including the SDP forms and interval positivity certificates, are given in Appendixes A and B.

A. Toy spectral density and input data

The toy spectral density is chosen to have a nontrivial matrix structure while remaining simple enough to evaluate exactly. For $a, b \in \{0, 1\}$, define

$$\rho_{ab}^{\text{toy}}(E) = \sum_{k=0}^{N_{\text{state}}-1} Z_{ka}^* Z_{kb} \delta(E - E_k), \quad (69)$$

with evenly spaced energies

$$E_k = \frac{1}{10}(k+1). \quad (70)$$

and $N_{\text{state}} = 96$. The overlap factors are

$$Z_{n0} = \begin{cases} +\frac{1}{5} & n=0, \\ +1 & n \text{ odd}, \\ -\frac{1}{5} & n > 0 \text{ even}, \end{cases} \quad Z_{n1} = \begin{cases} 1 & n \text{ even}, \\ \frac{1}{5} & n \text{ odd}. \end{cases} \quad (71)$$

These overlap factors are chosen to mirror a situation where the operator of interest has a relatively low overlap onto the ground state. This is a common source of difficulty in spectral applications, since the target observable may be sensitive to low-energy states that are not cleanly isolated by the target operator.

The Euclidean correlator generated by Eq. (69) is

$$C_{t,ab} = \sum_{k=0}^{N_{\text{state}}-1} Z_{ka}^* Z_{kb} e^{-E_k t}. \quad (72)$$

In all examples below, the correlator is sampled at $N_t = 20$ time slices. The scalar-correlator bounds use only the $C_{00}(t)$ component. The matrix-correlator bounds use the full 2×2 correlator matrix, while the target observable is still chosen to involve the same component of the spectral density as in the scalar case. This makes the scalar and matrix calculations directly comparable: the target is fixed, while the amount of positive-semidefinite input data is varied.

For moment-problem inputs, the data are the Euclidean correlator values in Eq. (72). For inputs to Nevanlinna-Pick-based interpolation, we instead use the Stieltjes transform of the same toy density,

$$G_{ab}(z) = \sum_{k=0}^{N_{\text{state}}-1} \frac{Z_{ka}^* Z_{kb}}{E_k - z}. \quad (73)$$

The transform is evaluated on the imaginary axis at $z = i\omega_n = 2\pi i n / N_t$ for $n = 0, \dots, N_t/2 - 1$. For numerical convenience, the point at zero frequency is shifted slightly to $\omega_0 = 10^{-3}\omega_1$, avoiding a point on the boundary of the upper half plane. The moment-problem and Nevanlinna-Pick-interpolation calculations therefore use different finite sets of input functionals of the same underlying toy density.

To model statistical uncertainty, the covariance matrix is taken to be

$$\Sigma_{tab,t'a'b'} = \alpha^2[s + (1-s)\delta_{tt'}\delta_{aa'}\delta_{bb'}]\Sigma_{tab,t'a'b'}^{(0)}, \quad (74)$$

where α is an overall scale parameter, $s = 1/2$ is a shrinkage parameter, and $\Sigma^{(0)}$ is defined by

$$\Sigma_{tab,t'a'b'}^{(0)} = \hat{C}_{tab}\hat{C}_{t'a'b'}e^{-\gamma|t-t'|}, \quad (75)$$

with $\gamma^{-1} = 1.3$. For each numeric example, $\hat{C}_t$ is drawn from a multivariate normal distribution $\hat{C} \sim \mathcal{N}(C, \Sigma)$.

For simplicity, the covariance matrix is taken to have the same structure in the time domain and in the upper half plane. The cutoff σ_0 [cf. Eq. (20)] is taken to be

$$\sigma_0 = \sqrt{2N_{\text{corr}}}\alpha, \quad (76)$$

which roughly corresponds to a cutoff of 2 on the χ^2 per “degree of freedom.”

B. Spectral reconstruction with Cauchy kernel

This section presents results for a Cauchy-smeared spectral reconstruction,

$$\tilde{\rho}_\epsilon(E) = \int d\lambda \tilde{K}_{E,\epsilon}^{\text{Cauchy}}(\lambda) \tilde{\rho}(\lambda), \quad (77)$$

$$\tilde{K}_{E,\epsilon}^{\text{Cauchy}}(\lambda) = \frac{\epsilon}{(\lambda - e^{-E})^2 + \epsilon^2}, \quad (78)$$

where $\tilde{\rho}(\lambda) = \lambda\rho(-\log \lambda)$ is the λ -space spectral function as in Eq. (9). Since $\tilde{K}_{E,\epsilon}^{\text{Cauchy}}$ is a rational function of λ , the problem of bounding $\rho_\epsilon(E)$ can be reformulated exactly as a finite-dimensional SDP via Eq. (52). By Eq. (14), the parameter ϵ corresponds geometrically to a fixed distance above the real line in the complex- λ plane.

Figure 1 shows the resulting bounds from solving Eq. (52). The reconstruction used the toy-model inputs for $\hat{C}_t$ (cf. Sec. VIA above) on $N_t = 20$ total time slices. A fixed value of $\epsilon = 0.1$ was used for the choice of the smearing kernel. As a stand-in for precise but realistic Monte Carlo data, the noise level of the input data is taken to be $\alpha = 10^{-4}$.

Figure 1 shows that the smeared spectral function $\tilde{\rho}_\epsilon(E)$ is bounded with good precision. As expected, the bounds from the matrix-valued input data are generally stronger than those from scalar-valued inputs.³

³The fact that the matrix bounds are not *always* better is a consequence of choosing to scale the cutoff σ_0 with N_{corr} , which is different in the scalar and matrix cases. For fixed cutoff σ_0 , the matrix bound will always be at least as tight as the scalar bound.




FIG. 1. Bounds on the smeared spectral function with kernel $\tilde{K}^{\text{Cauchy}}$ as defined in Eq. (78).

More generally, Fig. 1 demonstrates how these bounds can be used with noisy input data to constrain energy-dependent smeared spectral quantities. As the preceding sections show, the convex-optimization formulation offers the freedom to choose kernels besides Eq. (78).

C. Toy model for inclusive τ decay

Inclusive hadronic decays of the τ lepton provide a useful example of a threshold spectral observable. They are phenomenologically interesting because they can constrain Cabibbo–Kobayashi–Maskawa (CKM) matrix elements and because persistent tensions with CKM unitarity have motivated renewed lattice interest; see Refs. [57,58] and references therein. The inclusive hadronic decay rate can be written in terms of longitudinal and transverse spectral densities as [13,14]

$$\frac{R_{uq}^{(\tau)}}{|V_{uq}|^2} \propto \sum_{X=T,L} \int_0^\infty dE K_X\left(\frac{E}{m_\tau}\right) E^2 \rho_{X,uq}(E), \quad (79)$$

where $q = d, s$ and the omitted prefactors are known. The kernels are

$$K_L(x) = \frac{1}{x}(1-x^2)^2\theta(1-x), \quad (80)$$

$$K_T(x) = (1+2x^2)K_L(x). \quad (81)$$

The step function encodes the kinematic threshold at $E = m_\tau$.

For the numerical demonstration, we isolate a single channel and define

$$R_X = \int_0^\infty dE K_X\left(\frac{E}{m_\tau}\right) E^2 \rho_X(E), \quad X \in \{L, T\}. \quad (82)$$

Setting $E^2 \rho_X(E) = \rho^{\text{toy}}(E)$ allows computation of bounds on R_X . Numerical results will be presented below for R_L only with $m_\tau = 0.35$.

For analyticity-based interpolation, the kernel is a rational function of the energy variable on each side of the threshold, so the piecewise-PMP treatment of Sec. IV C can be applied directly. In the moment-problem formulation, a change of variables $\lambda = e^{-E}$ converts the spectral representation of the correlator to

$$C_X(t) = \int_0^1 d\lambda \tilde{\rho}_X(\lambda) \lambda^t, \quad (83)$$

and the target observable becomes

$$R_X = \int_0^1 d\lambda \tilde{K}_X(\lambda) \tilde{\rho}_X(\lambda). \quad (84)$$

The transformed kernel $\tilde{K}_X(\lambda)$ is not polynomial in λ , so the moment-problem calculation uses a rational approximation to the target kernel. For the numerical test, a total of $N_t = 20$ time slices are used. For simplicity, an order-20 Taylor expansion of $(1-\lambda)\tilde{K}_X(\lambda)$ around $\lambda = 1$ is used to approximate the target kernel. For the parameters used here, this approximation reproduces the exact toy value of $\mathcal{K}[\rho]$ to better than one part in 10^{11} . This approximation error is therefore negligible for the present demonstration. Note that the approximation order of the kernel may be increased arbitrarily irrespective of N_t .

The kernel has an infrared singularity in the energy variable. For the toy model, this is avoided by imposing a conservative lower edge of support, $m_{\text{gap}} = 0.05$, which is half the true spectral gap in Eq. (70). For Nevanlinna-Pick-based interpolation, the interval is restricted to $E \in [m_{\text{gap}}, \infty)$. For moment inputs the corresponding interval is $\lambda \in [0, e^{-m_{\text{gap}}}]$. This restriction is part of the definition of the interval I in Eq. (52).

Figure 2 shows the resulting bounds on R_L from solving the matrix ($r = 2$) versions of Eqs. (52) and (53) using moment-problem and NP-based inputs, respectively. In all cases, the bounds contain the exact value




FIG. 2. Bounds on the toy inclusive τ -decay observable R_L defined in Eq. (82). The central data are generated from the exact toy spectral density in Eq. (69), and the input uncertainty is varied through the modeled compatibility region. The difference between the sizes of the bounds is discussed in the main text.

computed directly from the toy spectral density. The bounds shrink as the input uncertainty is reduced, as expected from the decreasing size of the feasible set. For this particular observable and finite input data, the moment formulation happens to give substantially tighter bounds than the NP-based formulation. This should not be interpreted as an intrinsic hierarchy between the methods. Rather, the two calculations impose different finite-data constraints on the same underlying toy density.

D. Toy model for hadronic vacuum polarization

The next numerical demonstration models the calculation of the leading-order hadronic vacuum polarization (HVP) contribution to the anomalous magnetic moment of the muon using the toy model. This observable can be written in spectral form as

$$a_\mu^{\text{HVP}} = \left(\frac{\alpha_{\text{EM}}}{\pi}\right)^2 \int_0^\infty ds K(s) \rho(s), \quad (85)$$

where the kernel $K(s)$ is known analytically in terms of the squared energy s and is reviewed in Appendix E for

completeness. The corresponding Euclidean correlator admits the representation

$$C(t) = \frac{1}{2} \int_0^\infty ds \sqrt{s} \rho(s) e^{-\sqrt{s}t}. \quad (86)$$

Changing variables to $\lambda = e^{-\sqrt{s}}$ converts this into a Hausdorff moment problem,

$$C(t) = \int_0^1 d\lambda \tilde{\rho}(\lambda) \lambda^t, \quad (87)$$

with

$$\tilde{\rho}(\lambda) = \frac{1}{\lambda} (-\log \lambda)^2 \rho((-\log \lambda)^2). \quad (88)$$

In terms of this transformed density,

$$a_\mu^{\text{HVP}} = \int_0^1 d\lambda \tilde{\rho}(\lambda) \tilde{K}(\lambda), \quad (89)$$

where

$$\tilde{K}(\lambda) = \frac{2}{-\log \lambda} K((-\log \lambda)^2). \quad (90)$$

As in the preceding example, the transformed kernel is approximated before applying the PMP reduction. For simplicity, the numerical test uses a 20th-order Chebyshev approximation as implemented in MPMATH [56]. As above, the approximation error on the kernel was verified to be negligible and can be reduced arbitrarily. The HVP kernel is also singular at threshold, and the upper-bound calculation is stabilized by imposing the same explicit mass gap $m_{\text{gap}} = 0.05$. The purpose of the example is not to optimize the treatment of the physical HVP kernel, but to test how the bounds behave for a phenomenologically motivated kernel with nontrivial infrared structure.

Figure 3 shows the bounds resulting from solutions to the scalar ($r = 1$) and matrix ($r = 2$) versions of Eq. (52) using input data on $N_t = 20$ time slices. The exact value computed using the toy spectral function lies within the computed bounds throughout the noise range considered. As above, the cutoff σ_0 is scaled for the scalar and matrix cases according to Eq. (76). For larger input uncertainties (roughly the right-hand side of Fig. 3), the scalar and matrix problems give comparable bounds. As the input uncertainty is reduced, the bounds from the matrix problem become stronger, in agreement with expectations.

This behavior illustrates an important qualitative point. Adding operators is generally expected to improve spectral bounds, since the matrix positivity constraint contains additional information that is not present in a scalar correlator. On the other hand, a detailed statistical comparison necessarily depends both on the statistical




FIG. 3. Bounds on the toy HVP-like observable defined in Eq. (85). The scalar bounds use only the $C_{00}(t)$ component of the toy correlator, while the matrix bounds use the full 2×2 correlator matrix for the same target spectral component.

properties of the inputs (via the covariance matrix) and the choice of statistical cutoff (via σ_0). The choice of “equivalent” σ_0 values for different problem sizes must therefore be made with care. In this sense, matrix-valued correlators should be viewed as a promising source of additional constraints, not as an automatic improvement independent of statistics and the overall computational cost of generating the matrix of correlators.

VII. COMPARISON TO OTHER METHODS

This section compares the spectral-bootstrap construction with related reconstruction, analyticity, and optimization methods.

A. Kernel reconstruction methods

The dual formulation provides an interesting connection to methods that approximate the target smearing kernel as a linear combination of basis functions in order to infer $\mathcal{K}[\rho]$ from the linear combination of the same expansion coefficients with $\hat{C}_t$ [2,5,12]. For simplicity, consider a scalar spectral reconstruction problem with kernel $K: I \rightarrow \mathbb{R}$ and basis functions $b_t: I \rightarrow \mathbb{R}$. From a given approximation of the target smearing kernel

$$\bar{K}(x) = \sum_{t=1}^{N_t} g_t b_t(x), \quad (91)$$

the spectral observable smeared with the approximate $\bar{K}(x)$ is obtained directly from the correlator data:

$$\bar{\mathcal{K}}[\rho] = \int_I dx \rho(x) \bar{K}(x) = \sum_{t=1}^{N_t} g_t C_t. \quad (92)$$

A common approach is to choose coefficients g_t that approximate the target kernel, while controlling the noise amplification on $\bar{\mathcal{K}}[\rho]$.⁴ Often, a covariance penalty is introduced to tame large fluctuations in the coefficients g_t . We note that the noisy-data dual in Eq. (41) contains a similar term. While this term might be interpreted as yielding a similar effect, it bears emphasizing that the present bootstrap formulation chooses σ_0 as part of the definition of the statistical consistency region; σ_0 is not tuned as part of a bias-variance trade-off.

The dual bootstrap problem provides a similar kernel-reconstruction interpretation, but with two major advantages. The first is the fact that, should an approximation of the kernel be required for PMP reduction, the order of the approximation is conceptually unrelated to the number of available input data. This approximated kernel will be denoted $\underline{K}$ to distinguish it from $\bar{K}$ above. The same holds for approximating rational functions. We emphasize that the need to quantify the effects of such approximations, $\bar{K}$ and $\underline{K}$, is common to both approaches.

The second advantage is related to the actual role of N_t , which instead controls the approximation order of the *bounding* polynomial [cf. Eq. (49)]. The upshot is that *exact* inequalities replace the approximate equality above.

For a lower bound, the pointwise dual feasibility condition is

$$\underline{K}(x) \geq \sum_{t=1}^{N_t} g_t b_t(x) \quad \forall x \in I. \quad (93)$$

If $\rho(x) \geq 0$, this implies

$$\underline{K}[\rho] = \int_I dx \rho(x) \underline{K}(x) \geq \sum_{t=1}^{N_t} g_t C_t. \quad (94)$$

Thus, a feasible set of dual coefficients gives a one-sided reconstruction of the target observable. Upper bounds are obtained analogously. Clearly, the strength of the resulting rigorous bounds on $\mathcal{K}[\rho]$ necessarily scales with the quality of the bounding polynomials, and therefore with N_t .

The lattice bootstrap program presented in this work also offers a third advantage over current linear formulations of spectral reconstructions: it naturally extends to matrix-valued spectral functions. While state-of-the-art linear methods lack this extension, this work might serve as inspiration for future developments. Building on the analogies observed here between kernel reconstruction methods and the dual programs of this

work, a matrix-valued formulation of the linear methods appears natural. For instance, one could envision defining matrix-valued approximate kernels, as in Eq. (91), with matrix-valued coefficients g . When applied to the methods of Ref. [5], for instance, these coefficients could then be obtained from the minimization of a properly modified functional that incorporates the Hermitian trace pairing of the approximate and the target smearing kernels.

A notable class of smeared observables is constituted by those with an associated kernel whose linear expansion on the Euclidean time basis is exact. Within this class, time-momentum representations and window observables have played an important role, e.g., in determinations of hadronic contributions to the anomalous magnetic moment of the muon [61]. In the moment-problem language, a finite window observable can be defined in terms of a kernel of the form

$$K_{\text{win}}(\lambda) = \sum_{t=t_{\min}}^{t_{\max}} w_t \lambda^t, \quad (95)$$

or as the correspondingly discretized integral. If the correlator data are known exactly for all times appearing in the sum, then

$$\int_0^1 d\lambda K_{\text{win}}(\lambda) \tilde{\rho}(\lambda) = \sum_{t=t_{\min}}^{t_{\max}} w_t C_t \quad (96)$$

is obtained directly. In such cases, the dual certificate is exact with $g_t = w_t$, and there is no inverse problem to solve. However, Euclidean data are not always precisely available for all values of t , or the exact coefficients w_t might still yield unstable reconstructions of the smeared observable. This could be the case, e.g., for observables that require long-distance correlator information or for kernels with sharp spectral features. For these cases, casting the problem as a spectral reconstruction may possibly yield benefits, especially if matrix-valued positivity can be leveraged advantageously.

B. Comparison to the conformal bootstrap

The finite-dimensional reduction used in the present work is closely related to the numerical conformal bootstrap. Both approaches start from a positivity condition, formulate a polynomial matrix program, reduce it to a finite SDP, and solve the resulting problem with primal-dual interior point methods [37,52]. In this technical sense, the computational pipeline is the same.

The physical role of the input data is different. In the conformal bootstrap, positivity and crossing symmetry constrain possible conformal field theories without Monte Carlo data. In the present work, positivity is combined with finite Euclidean correlator data computed

⁴Note that Refs. [23,59,60] decouple the choice of such a regulator from the covariance matrix. Different choices of the regulator can be mapped to a corresponding measure of the distance of $C[\rho]$ from $\bar{C}$ in the definition of primal problems and might be worth studying in the future.

numerically, with no connection to conformal field theory. In distinction to the conformal bootstrap and to the older S -matrix bootstrap program, the present method does not attempt to solve QCD from general principles alone. Rather, convex optimization is used to propagate the results of nonperturbative calculations into rigorous bounds on smeared spectral observables.

This distinction also explains why two-point functions play a central role here, as opposed to the four-point functions typically used in the conformal bootstrap. In lattice QCD, two-point functions of chosen operators are nontrivial computable data. The positivity of their spectral representation already gives useful constraints on smeared spectral quantities. Higher-point functions could be incorporated in future generalizations, but they are not needed for the basic spectral-bounding problem studied in this work.

C. Alternative optimization algorithms

The dual problems can also be optimized without converting them to PMPs. For example, Ref. [29] considered scalar problems using interior-point or barrier methods applied directly to the dual. In that approach, the constrained optimization is replaced by an unconstrained or partially constrained optimization with a barrier term that discourages violation of dual feasibility.

Barrier and exchange methods have an important practical advantage. They do not require the polynomial or rational structure needed for the finite-dimensional PMP reduction. They can therefore be applied to a broader class of kernels and basis functions. They may also involve fewer variables than the finite SDP obtained from a sum-of-squares representation.

The trade-off is certification. If the positivity condition is enforced by sampling or discretizing the spectral domain, additional work is needed to ensure that no violation occurs between sample points. Likewise, if the method optimizes only the dual problem, the usual finite-SDP duality gap is not automatically available as a stopping certificate. These issues are not necessarily fatal, but they require separate control, for instance through adaptive refinement, interval arithmetic, or recovery of a matching primal feasible solution.

It is therefore not obvious *a priori* which numerical strategy will be best in practical lattice applications. The PMP-to-SDP route gives strong certificates for polynomial and rational problems, while barrier or exchange methods may be more flexible for exploratory work or for kernels not well suited to polynomial representation.

D. Other approaches to spectroscopy

Many lattice methods attack spectral problems by directly estimating finite-volume energies and overlap factors. Examples include generalized eigenvalue methods, multistate fits, Lanczos/Rayleigh-Ritz methods, and Prony

methods. These methods are closely related to moment problems and to the Krylov-space interpretation of Euclidean correlator data [62–66].

The goal of the present work is different. Here the primary output is not a direct calculation of the finite-volume energy spectrum, but instead a rigorous allowed interval for a smeared spectral observable. As discussed in Sec. III B, complementary slackness can identify spectra of extremal feasible measures, but those spectra should be interpreted as optimizers of the bounding problem rather than as estimates of the true physical spectrum.

The two perspectives are complementary. Spectrum-extraction methods can provide detailed physical interpretation and may give very efficient determinations when a small number of states dominate. Bootstrap bounds avoid committing to a particular finite-state fit ansatz and instead ask what can be concluded from positivity and the available correlator data alone. Which approach is more useful depends on the observable, the quality of the data, and the extent to which the relevant spectral region is resolved.

VIII. DISCUSSION AND OUTLOOK

This work has developed a spectral bootstrap framework for bounding smeared spectral observables from Euclidean lattice data. Starting from spectral positivity, the method formulates the problem as a convex optimization over all positive spectral densities compatible with the input data. The dual variables provide checkable one-sided kernel bounds, and polynomial, rational, or piecewise rational instances can be reduced to finite-dimensional semidefinite programs.

The remainder of this section summarizes the practical assumptions behind the construction, the uncertainties that are included, and the additional systematics that remain.

Interpretation of bounds for noisy data. Bounds for noisy data were discussed in Sec. II D. The main questions regarding interpretation are how the covariance matrix is estimated, whether it is regularized, and how the cutoff σ_0 is assigned. The interpretation also depends on the number of data components. Care must be taken, for example, when comparing scalar versus matrix bounds.

Additional uncertainty from kernel approximation. A separate source of uncertainty arises when the target kernel is approximated to obtain a finite-dimensional SDP. As discussed in Sec. IV, the finite reduction is exact for polynomial, rational, and piecewise rational kernels satisfying the stated denominator conditions. For an approximated kernel, the bootstrap bounds apply to the approximating observable; the difference from the desired observable is an additional systematic that must be bounded, monitored by convergence tests, or otherwise included in the final uncertainty budget.

Reflection positivity. The basic physical input used throughout this work is spectral positivity, which derives

from reflection positivity. Many large-scale lattice calculations use improved actions that break reflection positivity at finite lattice spacing in order to improve the approach to the continuum limit. For such actions, the transfer matrix may fail to be Hermitian at finite lattice spacing, and the earliest Euclidean time slices can contain unphysical oscillatory contributions. Applying the present bounds as exact finite-lattice-spacing statements then requires additional care. Possible strategies include omitting affected early-time data, inflating their uncertainties, or using a transfer matrix that advances by multiple lattice spacings when such a positive construction is available.

This issue is not unique to the present method. Many lattice analyses implicitly rely on reflection positivity even when it is not strictly present at finite lattice spacing. The appropriate level of rigor will depend on the intended application and the size of violations compared to other systematic uncertainties.

Matrix-valued correlators. Matrix correlators are particularly natural in this framework. Their positive-semidefinite spectral representation supplies constraints that are invisible in a single correlator. This is different from purely linear reconstructions, where additional operators do not affect the target observable unless they enter the chosen estimator. The toy HVP example in Sec. VID illustrates that matrix information can tighten the bounds substantially in favorable regimes.

Such improvement is not necessarily automatic. A larger correlator matrix also has more covariance components, may require a larger statistical cutoff at fixed confidence level, and may be more expensive to compute. Whether additional operators are worthwhile (e.g., compared to increasing statistics or extending the temporal dimension for fixed correlator basis) depends on their overlaps with the spectral region relevant to the target kernel, their noise properties, and the cost of computing the full matrix. Understanding this trade-off is an important practical question for future applications.

ACKNOWLEDGMENTS

We gratefully acknowledge useful conversations with Norman Christ and Sebastian Mizera. R. A. is an Ernest Kempton Adams Postdoctoral Fellow supported in part by the Ernest Kempton Adams fund for Physical Research of Columbia University. R. A. and S. F. are supported in part by U.S. DOE Grant No. DE-SC0011941. S. F. also is supported in part by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-2036197. P. O. is supported in part by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under Grant Contract No. DE-SC0012704 (BNL). The numerical calculations in this work made use of NumPy [67,68], FLINT [69], SymPy [70], SciPy [71], and MPMATH [56]. Figures were generated using Matplotlib [72].

Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

DATA AVAILABILITY

Code for the numeric examples in Sec. VI is available at [41].

APPENDIX A: FINITE-DIMENSIONAL SDP REDUCTION

This appendix collects the technical details used in Sec. IV to reduce polynomial matrix programs to finite-dimensional semidefinite programs. The construction follows the methods used in conformal-bootstrap applications of polynomial matrix programs [37], with minor modifications for the intervals and noisy-data constraints used here.

Following Ref. [37], these results are stated for symmetric matrix polynomials $M(x) \in \mathbb{S}^r[x]$, but as proven in Ref. [73] the same results hold for Hermitian matrix polynomials $M \in \mathcal{H}^r[x]$, provided that the certificate matrices Y_a are taken to be Hermitian rather than symmetric.

Let $q_0, q_1, \dots$ be a basis of real polynomials with $\deg q_i = i$. Define the matrix of polynomial products

$$Q_{ij}(x) = q_i(x)q_j(x). \quad (\text{A1})$$

The monomial basis $q_i(x) = x^i$ is the simplest conceptual choice, although orthogonal polynomial bases are typically preferred in numerical implementations.

For a polynomial matrix $M(x) \in \mathbb{S}^r[x]$, positivity on the real line can be represented by a matrix sum of squares. One convenient form is the following.

Lemma 3. Let $M(x) \in \mathbb{S}^r[x]$ have degree d . Then $M(x) \geq 0$ for all $x \in \mathbb{R}$ if and only if there exists a positive-semidefinite block matrix $Y \in \mathbb{S}^{[d/2]+1}$ such that

$$M(x) = \text{Tr}'[Y(Q(x) \otimes \mathbb{1}_r)], \quad (\text{A2})$$

where $\mathbb{1}_r$ is the $r \times r$ identity matrix and Tr' denotes the partial trace over the polynomial-basis indices, leaving an $r \times r$ matrix.

Equivalently, one may write

$$M(x) = V(x)^\top Y V(x), \quad (\text{A3})$$

where $V(x) = (q_0(x), q_1(x), \dots, q_m(x)) \otimes \mathbb{1}_r$ for $m = \lfloor d/2 \rfloor$.

For a half line, the corresponding representation is as follows.

Lemma 4. Let $M(x) \in \mathbb{S}^r[x]$ have degree d . Then $M(x) \geq 0$ for all $x \geq 0$ if and only if there exist

positive-semidefinite block matrices $Y_1 \in \mathbb{S}^{[d/2]+1}$ and $Y_2 \in \mathbb{S}^{[(d-1)/2]+1}$ such that

$$M(x) = \text{Tr}'[Y_1(Q(x) \otimes \mathbb{1}_r)] + x \text{Tr}'[Y_2(Q(x) \otimes \mathbb{1}_r)]. \quad (\text{A4})$$

By translating the variable, the same representation holds for any half line $[a, \infty)$ with the lemma applied to $M(a+x)$ with $x \geq 0$.

For bounded intervals, one may use the standard Hausdorff-moment positivity certificates. For the interval $[0, 1]$, the result can be written as follows [73].

Lemma 5. (Dette-Studden). Let $M(x) \in \mathbb{S}^r[x]$ have degree d . If d is even, then $M(x) \geq 0$ for all $x \in [0, 1]$ if and only if there exist positive-semidefinite block matrices $Y_1 \in \mathbb{S}^{d/2+1}$ and $Y_2 \in \mathbb{S}^{d/2}$ such that

$$M(x) = \text{Tr}'[Y_1(Q(x) \otimes \mathbb{1}_r)] + x(1-x) \text{Tr}'[Y_2(Q(x) \otimes \mathbb{1}_r)]. \quad (\text{A5})$$

If d is odd, then $M(x) \geq 0$ for all $x \in [0, 1]$ if and only if there exist positive-semidefinite block matrices $Y_1 \in \mathbb{S}^{[d/2]+1}$ and $Y_2 \in \mathbb{S}^{[d/2]+1}$ such that

$$M(x) = x \text{Tr}'[Y_1(Q(x) \otimes \mathbb{1}_r)] + (1-x) \text{Tr}'[Y_2(Q(x) \otimes \mathbb{1}_r)]. \quad (\text{A6})$$

The corresponding certificate on a general bounded interval $[a, b]$ is obtained by an affine change of variables mapping $[a, b]$ to $[0, 1]$.

APPENDIX B: PRIMAL-DUAL INTERIOR POINT METHODS

This appendix reviews the solution of finite-dimensional SDPs like those in Sec. IV using primal-dual interior point methods. Appendix B 1 begins by defining a dual version of the finite SDP in Eq. (58) and defining the duality gap. Appendix B 2 then describes their solution using primal-dual interior point methods; Appendix B 2 follows closely the presentation of Ref. [37], with a few modifications specific to the SDPs encountered in this work.

1. Primal-dual finite SDP and the duality gap

Consider a generic finite SDP of the form

$$\begin{aligned} & \underset{y, Y}{\text{maximize}} && b \cdot y + D \cdot Y \\ & \text{subject to} && \text{Tr}[A_p Y] + (By)_p = c_p, \\ & && \forall p \in \{1, \dots, P\}, \\ & && Y \succeq 0, \end{aligned} \quad (\text{B1})$$

for which the Lagrange dual has the form

$$\begin{aligned} & \underset{x}{\text{minimize}} && c \cdot x \\ & \text{subject to} && X = \sum_p A_p x_p - D, \\ & && B^\top x = b, \\ & && X \succeq 0. \end{aligned} \quad (\text{B2})$$

This finite-dimensional dual pair should not be confused with the original primal and dual spectral reconstruction problems. The terminology is inherited from the SDP solver: Eq. (B2) is the primal finite SDP, while Eq. (B1) is the dual finite SDP.

If (x, X) and (y, Y) are feasible for the finite primal and dual SDPs, then

$$c \cdot x - D \cdot Y - b \cdot y = X \cdot Y \geq 0. \quad (\text{B3})$$

The left-hand side is the finite-SDP duality gap. Primal-dual interior point methods monitor this quantity during optimization. A sufficiently small duality gap, together with feasibility of the finite SDP variables, certifies that the finite SDP has been solved to the corresponding precision.

Further implementation details of the primal-dual interior point method used in the numerical examples are reviewed in Appendix B.

2. Interior point methods

The starting point for primal-dual interior point methods is the finite SDP dual pair in Eqs. (B1) and (B2). For feasible primal and dual SDP variables, the duality gap is $X \cdot Y \geq 0$. Since $X, Y \succeq 0$, this gap vanishes if and only if the matrix product XY vanishes.

Primal-dual methods replace the exact complementarity condition by the deformed equation

$$XY = \mu \mathbb{1}_{\dim X}, \quad (\text{B4})$$

where $\mu > 0$ and $\mathbb{1}_{\dim X}$ is the identity matrix of the same dimension as X . Together with primal and dual feasibility, Eq. (B4) defines the central path

$$(x(\mu), X(\mu), y(\mu), Y(\mu)). \quad (\text{B5})$$

The optimization proceeds by Newton steps that push the variables toward the central path while gradually taking $\mu \rightarrow 0$.

To derive the Newton step, expand

$$(x, X, y, Y) \mapsto (x + dx, X + dX, y + dy, Y + dY) \quad (\text{B6})$$

and solve the linearized central-path equations. It is useful to define the residuals

$$P = \sum_p A_p x_p - X - D, \quad (\text{B7})$$

$$p = b - B^\top x, \quad (\text{B8})$$

$$d_i = c_i - \text{Tr}[A_i Y] - (By)_i, \quad (\text{B9})$$

$$R = \mu I - XY. \quad (\text{B10})$$

In terms of these residuals, the vector steps dx and dy are obtained from

$$\begin{pmatrix} S & -B \\ B^\top & 0 \end{pmatrix} \begin{pmatrix} dx \\ dy \end{pmatrix} = \begin{pmatrix} -d_i - \text{Tr}[A_i Z] \\ p \end{pmatrix}, \quad (\text{B11})$$

where $Z = X^{-1}(PY - R)$, and S is referred to as the Schur complement and is defined by

$$S_{ij} = \text{Tr}[A_i X^{-1} A_j Y]. \quad (\text{B12})$$

For the rank-one choices of A_i used in this work, the trace in Eq. (B12) can be evaluated efficiently by changing the order of index contraction, as in Ref. [37]. Once dx is known, the matrix steps are

$$dX = P + \sum_i A_i dx_i, \quad (\text{B13})$$

$$dY = X^{-1}(R - dXY). \quad (\text{B14})$$

The expression for dY is not necessarily symmetric. Following SDPB, we symmetrize dY before taking the step.

There are several implementation details in which the solver used here differs from the original implementation of Ref. [37]. First, the implementation allows only a subset of the linear equalities to be enforced, as needed for the noisy-data SDP constraints in Eq. (36). Second, following later versions of SDPB, we do not add the identity matrix to the left-hand side of Eq. (B11) [52]. Third, we use the same step size for the primal and dual variables, taking the smaller allowed value when necessary. Occasional stalling is handled by increasing μ by a factor of 1.5 and holding it fixed for ten Newton iterations. All numerical examples in this work use 150 decimal digits of precision.

The final modification is the elimination of free variables. This reduction was suggested in Ref. [37], following Ref. [74], but was not used there because of the structure of conformal-bootstrap SDPs. In the problems considered here, eliminating the free variables reduces the problem size and improves numerical stability.

The essential idea is to split the primal vector as

$$x = x_{\parallel} + x_{\perp}, \quad (\text{B15})$$

such that

$$B^\top x_{\parallel} = b, \quad B^\top x_{\perp} = 0. \quad (\text{B16})$$

The variables $x_{\parallel}$ can then be eliminated, which also eliminates the free variables y . Rather than constructing the relevant subspaces from linearly independent columns of B , we use an SVD factorization

$$B = U \text{diag}(\sigma_i) \tilde{V}^\top = UV^\top, \quad (\text{B17})$$

where the singular values have been absorbed into $V^\top$. The decomposition is arranged so that V is square and invertible, and

$$U = (U_{\parallel} \quad U_{\perp}), \quad (\text{B18})$$

with $U_{\parallel}$ square of the same dimension as V .

With this decomposition, the substitutions

$$\tilde{x} = U_{\parallel} x, \quad (\text{B19})$$

$$\tilde{c} = U_{\perp} c, \quad (\text{B20})$$

$$\tilde{D} = \sum_p A_p U_{\parallel, pq} (V^{-\top} b)_q, \quad (\text{B21})$$

$$\alpha_0 = c^\top U_{\parallel} V^{-\top} b, \quad (\text{B22})$$

$$\tilde{A}_q = \sum_p (U_{\perp})_{pq} A_p \quad (\text{B23})$$

reduce the original primal finite SDP to

$$\begin{aligned} &\text{minimize } \alpha_0 + \tilde{c} \cdot \tilde{x} \\ &\text{subject to } X = \sum_q \tilde{A}_q \tilde{x}_q - \tilde{D}, \\ &X \geq 0. \end{aligned} \quad (\text{B24})$$

Dualizing gives the reduced dual SDP

$$\begin{aligned} &\text{maximize } \tilde{D} \cdot Y \\ &\text{subject to } \text{Tr}[\tilde{A}_p Y] = \tilde{c}_p, \quad \forall p. \end{aligned} \quad (\text{B25})$$

The Schur complement for the reduced problem can be computed from the unreduced Schur complement by

$$\tilde{S}(X, Y) = U_{\perp}^\top S(X, Y) U_{\perp}. \quad (\text{B26})$$

This projection breaks some of the original block structure of S , as noted in Ref. [37]. For the SDPs considered here, the reduced dimension and improved conditioning outweigh this cost.

APPENDIX C: EQUIVALENT DESCRIPTIONS OF THE WERTEVORRAT

This appendix contains a proof of Lemma 1 and numerical demonstrations of both Lemmas 1 and 2. Since both sides of the equivalence can always be computed explicitly, the proof focuses on structural aspects, ignoring inessential details including measure-theoretic details of convergence and regularity and existence of relevant minimizations.

Proof of Lemma 1. Suppose that the moments $\hat{C}_t$ are given for a scalar nonextremal truncated Hamburger moment problem with $N = 2n$. Let $\mathcal{F} = \{\rho \in \mathcal{M}_+^1(\mathbb{R}) : C_t[\rho] = \hat{C}_t, t = 0, \dots, N\}$ be the set of feasible spectral functions. For z in the upper half plane, define the set of possible values of Green's function in Eq. (12) by

$$\Delta_N(z) = \text{cl}\{G_\rho(z) : \rho \in \mathcal{F}\}. \quad (\text{C1})$$

Kovalishina [53] has proven that $\Delta_N(z)$ is a disk

$$\Delta_N(z) = \bar{\mathbb{D}}(c, r) = \{\zeta \in \mathbb{C} : |\zeta - c| \leq r\}, \quad (\text{C2})$$

with known analytic formulas for the center c and radius r in terms of the moments $\hat{C}_t$.

On the other hand, the SDP reconstruction is obtained from the following intersection of half planes:

$$\mathbb{D}' = \bigcap_{\theta \in [0, 2\pi)} \{\zeta \in \mathbb{C} : \text{Re}(e^{i\theta}\zeta) \geq \ell(\theta)\}, \quad (\text{C3})$$

where, for each $\theta \in [0, 2\pi)$, we define the supporting lower bound

$$\ell(\theta) = \inf_{\rho \in \mathcal{F}} \text{Re}(e^{i\theta}G_\rho(z)). \quad (\text{C4})$$

Equivalently, assuming the relevant extremal values are attained after taking the closure and using Eq. (C1), we obtain

$$\ell(\theta) = \min_{\zeta \in \Delta_N(z)} \text{Re}(e^{i\theta}\zeta). \quad (\text{C5})$$

For each θ , the minimum of $\text{Re}(e^{i\theta}\zeta)$ is attained at the boundary point in the direction opposite to $e^{i\theta}$, and hence,

$$\ell(\theta) = \min_{\zeta \in \bar{\mathbb{D}}(c, r)} \text{Re}(e^{i\theta}\zeta) = \text{Re}(e^{i\theta}c) - r. \quad (\text{C6})$$

Thus, the SDP half-plane reconstruction is

$$\mathbb{D}' = \bigcap_{\theta \in [0, 2\pi)} \{\zeta \in \mathbb{C} : \text{Re}(e^{i\theta}(\zeta - c)) \geq -r\}. \quad (\text{C7})$$

If $\zeta \in \bar{\mathbb{D}}(c, r)$, then

$$\text{Re}(e^{i\theta}(\zeta - c)) \geq -|\zeta - c| \geq -r \quad \forall \theta, \quad (\text{C8})$$

so $\zeta \in \mathbb{D}'$, i.e., $\bar{\mathbb{D}}(c, r) \subseteq \mathbb{D}'$. Conversely, suppose $\zeta \in \mathbb{D}'$. If $\zeta = c$, then $\zeta \in \bar{\mathbb{D}}(c, r)$ trivially. Suppose therefore $x = \zeta - c \neq 0$. Choosing $\theta = \pi - \arg x$ gives

$$-|x| = \text{Re}(e^{i\theta}x) \geq -r. \quad (\text{C9})$$

Therefore, $|\zeta - c| = |x| \leq r$, so $\zeta \in \bar{\mathbb{D}}(c, r)$, i.e., $\mathbb{D}' \subseteq \bar{\mathbb{D}}(c, r)$. Hence,

$$\mathbb{D}' = \bar{\mathbb{D}}(c, r), \quad (\text{C10})$$

which was to be shown. ■

Figure 4 confirms the equivalence of the bounds for $G(z)$ from Eq. (12) computed using Lemma 1 and Lemma 2 at fixed locations in the upper half plane, respectively




FIG. 4. Numerical demonstration of the equivalent formulation of bounds on $G(z)$ in Eq. (12). The left panel shows the two different methods in Lemma 1 related to the moment problem. The right panel shows the two different methods related to Nevanlinna-Pick interpolation in Lemma 2.

$z = E + i\epsilon$ and $z = \lambda + i\epsilon$ at $E = 0.4$, $\epsilon = 0.1$ with $\lambda = e^{-E}$. For simplicity, the numerical example used the toy model from Sec. VI A with $N_{\text{states}} = 20$ and exact data on $N_t = 8$ time slices.

APPENDIX D: NEVANLINNA-PICK DATA FROM MOMENT DATA

This appendix records a direct relation between finite moment data and input data for analyticity-based interpolation methods like Nevanlinna-Pick interpolation. The relation is useful for comparing the two formulations, but it also highlights an important practical point: finite moment data and independently supplied values of a Stieltjes transform are not the same input unless the transformation between them is explicitly performed.

Start from a Hausdorff moment problem on $[0, 1)$,

$$C_t[\rho] = \int_0^1 d\lambda \rho(\lambda) \lambda^t, \quad t = 0, \dots, N_t - 1. \quad (\text{D1})$$

Let $\omega_n = 2\pi n/N_t$ for bosonic frequencies, or $\omega_n = 2\pi(n + \frac{1}{2})/N_t$ for fermionic frequencies. Define the discrete Fourier transform of the moment data by

$$C_{\omega_n}[\rho] = \sum_{t=0}^{N_t-1} e^{-i\omega_n t} C_t[\rho]. \quad (\text{D2})$$

Substituting Eq. (D1) into Eq. (D2) gives

$$\begin{aligned} C_{\omega_n}[\rho] &= \int_0^1 d\lambda \rho(\lambda) \sum_{t=0}^{N_t-1} (\lambda e^{-i\omega_n})^t \\ &= \int_0^1 d\lambda \rho(\lambda) \frac{1 - (\lambda e^{-i\omega_n})^{N_t}}{1 - \lambda e^{-i\omega_n}} \\ &= e^{i\omega_n} \int_0^1 d\lambda \rho(\lambda) \frac{1 \mp \lambda^{N_t}}{e^{i\omega_n} - \lambda}. \end{aligned} \quad (\text{D3})$$

The upper or lower sign is determined by whether $e^{i\omega_n N_t} = +1$ or -1 .

Define

$$z_n = e^{i\omega_n} \quad (\text{D4})$$

and introduce the modified non-negative density

$$\tilde{\rho}(\lambda) = (1 \mp \lambda^{N_t})\rho(\lambda), \quad (\text{D5})$$

with the same sign convention as in Eq. (D3). Then

$$C_{\omega_n}[\rho] = z_n \int_0^1 d\lambda \frac{\tilde{\rho}(\lambda)}{z_n - \lambda}. \quad (\text{D6})$$

Up to the overall factor of z_n , this is the Stieltjes transform of $\tilde{\rho}$ evaluated at z_n . With the convention

$$\tilde{G}(z) = \int_0^1 d\lambda \frac{\tilde{\rho}(\lambda)}{z - \lambda}, \quad (\text{D7})$$

the Fourier-transformed moment data specify

$$C_{\omega_n}[\rho] = z_n \tilde{G}(z_n). \quad (\text{D8})$$

There is redundancy in these interpolation conditions. For real moment data, $\tilde{G}(z^*) = \tilde{G}(z)^*$ and $z_n^* = z_{N_t-n}$. A nonredundant set of conditions can therefore be taken from the nodes in the upper half of the unit circle. In this sense, finite moment data can be repackaged as interpolation data for the Stieltjes transform on the unit circle.

The present formulation of the Nevanlinna-Pick interpolation problem on $\tilde{G}$ is directly accessible from Euclidean time correlator data. In contrast to its continuum counterpart in the energy variable E (cf. Sec. VI in Ref. [24]), the interpolation is phrased exactly in the upper half plane of the complex variable λ .

This construction explains why the two approaches can be compared within the same convex-optimization framework. It also explains why bounds obtained from independently chosen moment and Nevanlinna-Pick data need not coincide with finite data. They agree only when the two sets of inputs encode the same linear information about the same positive density and the same target functional.

APPENDIX E: HVP KERNEL

This appendix records the kernel used for the hadronic-vacuum-polarization example in Sec. VI D. The leading-order hadronic vacuum polarization contribution to the anomalous magnetic moment of the muon can be written as [61,75,76]

$$a_\mu^{\text{HVP}} = \left(\frac{\alpha}{\pi}\right)^2 \int_0^\infty dQ^2 K_E(Q^2) \hat{\Pi}(Q^2), \quad (\text{E1})$$

where

$$K_E(s) = \frac{1}{m_\mu^2} \hat{s} Z(\hat{s})^3 \frac{1 - \hat{s} Z(\hat{s})}{1 + \hat{s} Z(\hat{s})^2}, \quad (\text{E2})$$

$$Z(\hat{s}) = \frac{\hat{s} - \sqrt{\hat{s}^2 + 4\hat{s}}}{2\hat{s}}, \quad \hat{s} = \frac{s}{m_\mu^2}. \quad (\text{E3})$$

Here m_μ is the muon mass and

$$\hat{\Pi}(Q^2) = 4\pi^2 [\Pi(Q^2) - \Pi(0)] \quad (\text{E4})$$

is the subtracted vacuum polarization, restricted to the hadronic contribution.

The electromagnetic vector-current correlator admits the spectral representation

$$C(t) = \frac{1}{2} \int_0^\infty ds \sqrt{s} \rho(s) e^{-\sqrt{s}t}, \quad (\text{E5})$$

where the positive spectral density $\rho(s)$ is proportional to the usual R -ratio spectral function. The same spectral density enters the dispersion relation

$$\Pi(Q^2) - \Pi(0) = Q^2 \int_0^\infty \frac{ds \rho(s)}{s(s + Q^2)}. \quad (\text{E6})$$

Substituting Eq. (E6) into Eq. (E1) gives the spectral representation

$$a_\mu^{\text{HVP}} = \left(\frac{\alpha}{\pi}\right)^2 \int_0^\infty ds K(s) \rho(s), \quad (\text{E7})$$

with

Let

$$\tau = \frac{s}{4m_\mu^2}, \quad (\text{E9})$$

and

$$x = \frac{1 - \beta_\mu}{1 + \beta_\mu}, \quad (\text{E10})$$

$$\beta_\mu = \sqrt{1 - \frac{4m_\mu^2}{s}}. \quad (\text{E11})$$

The numerical examples in Sec. VID take $m_\mu = 0.2$. The kernel can be evaluated in closed form in terms of these variables, taking the form [76–80]

$$K(s) = \begin{cases} \frac{1}{2} - 4\tau - 4\tau(1 - 2\tau) \log(4\tau) - 2(1 - 8\tau + 8\tau^2) \sqrt{\frac{\tau}{1-\tau}} \cos^{-1}(\sqrt{\tau}), & 0 \leq s \leq 4m_\mu^2, \\ \frac{x^2}{2} (2 - x^2) + \frac{(1+x^2)(1+x)^2}{x^2} \left[\log(1+x) - x + \frac{x^2}{2} \right] + \frac{1+x}{1-x} x^2 \log x, & s \geq 4m_\mu^2. \end{cases} \quad (\text{E12})$$

-
- [1] E. C. Poggio, H. R. Quinn, and S. Weinberg, Smearing the quark model, *Phys. Rev. D* **13**, 1958 (1976).
 - [2] J. C. A. Barata and K. Fredenhagen, Particle scattering in Euclidean lattice field theories, *Commun. Math. Phys.* **138**, 507 (1991).
 - [3] M. T. Hansen, H. B. Meyer, and D. Robaina, From deep inelastic scattering to heavy-flavor semileptonic decays: Total rates into multihadron final states from lattice QCD, *Phys. Rev. D* **96**, 094513 (2017).
 - [4] J. Bulava and M. T. Hansen, Scattering amplitudes from finite-volume spectral functions, *Phys. Rev. D* **100**, 034521 (2019).
 - [5] M. Hansen, A. Lupo, and N. Tantalo, Extraction of spectral densities from lattice correlators, *Phys. Rev. D* **99**, 094508 (2019).
 - [6] A. Patella and N. Tantalo, Scattering amplitudes from Euclidean correlators: Haag-Ruelle theory and approximation formulae, *J. High Energy Phys.* **01** (2024) 091.
 - [7] A. Jüttner, Standard Model tests with smeared experiment and theory, [arXiv:2603.15487](https://arxiv.org/abs/2603.15487).
 - [8] S. Jeon and L. G. Yaffe, From quantum field theory to hydrodynamics: Transport coefficients and effective kinetic theory, *Phys. Rev. D* **53**, 5799 (1996).
 - [9] H. B. Meyer, A Calculation of the shear viscosity in SU(3) gluodynamics, *Phys. Rev. D* **76**, 101701 (2007).
 - [10] H. B. Meyer, Transport properties of the quark-gluon plasma: A lattice QCD perspective, *Eur. Phys. J. A* **47**, 86 (2011).
 - [11] G. Bailas, S. Hashimoto, and T. Ishikawa, Reconstruction of smeared spectral function from Euclidean correlation functions, *Prog. Theor. Exp. Phys.* **2020**, 043B07 (2020).
 - [12] P. Gambino and S. Hashimoto, Inclusive semileptonic decays from lattice QCD, *Phys. Rev. Lett.* **125**, 032001 (2020).
 - [13] A. Evangelista, R. Frezzotti, N. Tantalo, G. Gagliardi, F. Sanfilippo, S. Simula, and V. Lubicz (Extended Twisted Mass Collaboration), Inclusive hadronic decay rate of the τ lepton from lattice QCD, *Phys. Rev. D* **108**, 074513 (2023).
 - [14] C. Alexandrou *et al.* (Extended Twisted Mass Collaboration), Inclusive hadronic decay rate of the τ lepton from lattice QCD: The $\bar{u}s$ flavor channel and the Cabibbo angle, *Phys. Rev. Lett.* **132**, 261901 (2024).
 - [15] G. Backus and F. Gilbert, The resolving power of gross earth data, *Geophys. J. Int.* **16**, 169 (1968).
 - [16] F. Pijpers and M. Thompson, Faster formulations of the optimally localized averages method for helioseismic inversions, *Astron. Astrophys.* **262**, L33 (1992), <https://adsabs.harvard.edu/full/1992A%26A...262L..33P>.

- [17] H. Fukaya, S. Hashimoto, T. Kaneko, and H. Ohki, Towards fully nonperturbative computations of inelastic ℓN scattering cross sections from lattice QCD, *Phys. Rev. D* **102**, 114516 (2020).
- [18] A. P. Valentine and M. Sambridge, Gaussian process models—I. A framework for probabilistic continuous inverse theory, *Geophys. J. Int.* **220**, 1632 (2019).
- [19] A. P. Valentine and M. Sambridge, Gaussian process models—II. Lessons for discrete inversion, *Geophys. J. Int.* **220**, 1648 (2019).
- [20] J. Horak, J. M. Pawłowski, J. Rodríguez-Quintero, J. Turnwald, J. M. Urban, N. Wink, and S. Zafeiropoulos, Reconstructing QCD spectral functions with Gaussian processes, *Phys. Rev. D* **105**, 036014 (2022).
- [21] A. Candido, L. Del Debbio, T. Giani, and G. Petrillo, Bayesian inference with Gaussian processes for the determination of parton distribution functions, *Eur. Phys. J. C* **84**, 716 (2024).
- [22] L. Del Debbio, A. Lupo, M. Panero, and N. Tantalo, Bayesian solution to the inverse problem and its relation to Backus–Gilbert methods, *Eur. Phys. J. C* **85**, 185 (2025).
- [23] M. Bruno, L. Giusti, and M. Saccardi, Spectral densities from Euclidean lattice correlators via the Mellin transform, *Phys. Rev. D* **111**, 094515 (2025).
- [24] T. Bergamaschi, W. I. Jay, and P. R. Oare, Hadronic structure, conformal maps, and analytic continuation, *Phys. Rev. D* **108**, 074516 (2023).
- [25] J. Fei, C.-N. Yeh, and E. Gull, Nevanlinna analytical continuation, *Phys. Rev. Lett.* **126**, 056402 (2021).
- [26] R. Abbott, W. I. Jay, and P. R. Oare, Moment problems and bounds for matrix-valued smeared spectral functions, *arXiv:2508.01377*.
- [27] S. Fields and N. Christ, Nevanlinna-Pick interpolation from uncertain data, *Phys. Rev. D* **113**, 054502 (2026).
- [28] R. Abbott, W. Jay, and P. Oare, Moment problems and spectral functions, *arXiv:2602.11260*.
- [29] S. Lawrence, Model-free spectral reconstruction via Lagrange duality, *arXiv:2408.11766*.
- [30] D. Simmons-Duffin, The Conformal Bootstrap, in *Theoretical Advanced Study Institute in Elementary Particle Physics: New Frontiers in Fields and Strings* (World Scientific, Boulder, Colorado, 2017), pp. 1–74, *arXiv:1602.07982*, 10.1142/10270.
- [31] X. Han, S. A. Hartnoll, and J. Kruthoff, Bootstrapping matrix quantum mechanics, *Phys. Rev. Lett.* **125**, 041601 (2020).
- [32] H. W. Lin and Z. Zheng, Bootstrapping ground state correlators in matrix theory. Part I, *J. High Energy Phys.* **01** (2025) 190.
- [33] M. Cho, B. Gabai, H. W. Lin, J. Yeh, and Z. Zheng, Bootstrapping Euclidean two-point correlators, *J. High Energy Phys.* **08** (2026) 147.
- [34] S. Lawrence, B. McPeak, and D. Neill, Bootstrapping time-evolution in quantum mechanics, *arXiv:2412.08721*.
- [35] M. Luscher and U. Wolff, How to calculate the elastic scattering matrix in two-dimensional quantum field theories by numerical simulation, *Nucl. Phys.* **B339**, 222 (1990).
- [36] B. Blossier, M. Della Morte, G. von Hippel, T. Mendes, and R. Sommer, On the generalized eigenvalue method for energies and matrix elements in lattice field theory, *J. High Energy Phys.* **04** (2009) 094.
- [37] D. Simmons-Duffin, A Semidefinite Program solver for the conformal bootstrap, *J. High Energy Phys.* **06** (2015) 174.
- [38] K. Isii, Inequalities of the types of chebyshev and cramér-rao and mathematical programming, *Ann. Inst. Stat. Math.* **16**, 277 (1964).
- [39] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, England, 2004).
- [40] S. Arora, B. Barak *et al.*, *Computational Complexity: A Modern Approach* (Cambridge University Press, Cambridge, England, 2009), Vol. 1.
- [41] R. Abbott *et al.*, Lattice bootstrap, https://github.com/rabbott999/lattice_bootstrap (2026).
- [42] D. Bertsimas and I. Popescu, Optimal inequalities in probability theory: A convex optimization approach, *SIAM J. Optim.* **15**, 780 (2005).
- [43] J. B. Lasserre, A semidefinite programming approach to the generalized problem of moments, *Math. Program.* **112**, 65 (2008).
- [44] N. Christ, R. Eranki, and C. Kelly (RBC, UKQCD Collaborations), Bootstrap-determined p values in lattice QCD, *Phys. Rev. D* **111**, 074514 (2025).
- [45] M. Bruno and R. Sommer, On fits to correlated and auto-correlated data, *Comput. Phys. Commun.* **285**, 108643 (2023).
- [46] M. Slater, Lagrange multipliers revisited, in *Traces and Emergence of Nonlinear Programming* (Springer, New York, 2013), pp. 293–306.
- [47] D. Simmons-Duffin, The lightcone bootstrap and the spectrum of the 3d Ising CFT, *J. High Energy Phys.* **03** (2017) 086.
- [48] S. El-Showk and M. F. Paulos, Bootstrapping conformal field theories with the extremal functional method, *Phys. Rev. Lett.* **111**, 241601 (2013).
- [49] C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems* (SIAM, Philadelphia, PA, 1995), 10.1137/1.9781611971217.
- [50] Z.-Q. Luo, J. F. Sturm, S. Zhang *et al.*, Duality and self-duality for conic convex programming, Technical Report EI 9620-/A Tinbergen Institute, 1996.
- [51] D. Hilbert, Über die darstellung definiten formen als summe von formenquadraten, *Math. Ann.* **32**, 342 (1888).
- [52] W. Landry and D. Simmons-Duffin, Scaling the semidefinite program solver SDPB, *arXiv:1909.09745*.
- [53] I. V. Kovalishina, Analytic theory of a class of interpolation problems, *Math. USSR-Izv.* **22**, 419 (1984).
- [54] M. G. Krein and A. A. Nudelman, An interpolation problem in the class of Stieltjes functions and its connection with other problems, *Integr. Equ. Oper. Theory* **30**, 251 (1998).
- [55] M. Heide, B. Fritzsche, B. Kirstein *et al.*, Weyl sets in a non-degenerate truncated matricial Hausdorff moment problem, *Complex Anal. Oper. Theory* **18**, 82 (2024).
- [56] The MPMATH development team, MPMATH: a Python library for arbitrary-precision floating-point arithmetic (version 1.3.0), <http://mpmath.org/> (2023).
- [57] S. Banerjee *et al.* (Heavy Flavor Averaging Group (HFLAV) Collaboration), Averages of b-hadron, c-hadron, and τ -lepton properties as of 2023, *Phys. Rev. D* **113**, 012008 (2026).
- [58] Y. Aoki *et al.* (Flavour Lattice Averaging Group (FLAG) Collaboration), FLAG review 2024, *Phys. Rev. D* **113**, 014508 (2026).

- [59] M. Bruno, L. Giusti, and M. Saccardi, On the prediction of spectral densities from Lattice QCD, *Nuovo Cimento C* **47**, 197 (2024).
- [60] A. Lupo and N. Tantalo, Extraction of spectral densities from lattice correlators: Decoupling signal from noise, [arXiv:2605.14652](https://arxiv.org/abs/2605.14652).
- [61] D. Bernecker and H. B. Meyer, Vector correlators in lattice QCD: Methods and applications, *Eur. Phys. J. A* **47**, 148 (2011).
- [62] M. Fischer, B. Kostrzewa, J. Ostmeyer, K. Ottnad, M. Ueding, and C. Urbach, On the generalised eigenvalue method and its relation to Prony and generalised pencil of function methods, *Eur. Phys. J. A* **56**, 206 (2020).
- [63] M. L. Wagman, Lanczos algorithm, the transfer matrix, and the signal-to-noise problem, *Phys. Rev. Lett.* **134**, 241901 (2025).
- [64] D. C. Hackett and M. L. Wagman, Block Lanczos for lattice QCD spectroscopy and matrix elements, *Phys. Rev. D* **112**, 014514 (2025).
- [65] J. Ostmeyer, A. Sen, and C. Urbach, On the equivalence of Prony and Lanczos methods for Euclidean correlation functions, *Eur. Phys. J. A* **61**, 26 (2025).
- [66] R. Abbott, D. C. Hackett, G. T. Fleming, D. A. Pefkou, and M. L. Wagman, Filtered Rayleigh–Ritz is all you need, *Phys. Rev. E* **114**, 015302 (2026).
- [67] S. van der Walt, S. C. Colbert, and G. Varoquaux, The NumPy array: A structure for efficient numerical computation, *Comput. Sci. Eng.* **13**, 22 (2011).
- [68] C. R. Harris *et al.*, Array programming with NumPy, *Nature (London)* **585**, 357 (2020).
- [69] T. F. team, FLINT: Fast library for number theory, version 3.5.0, <https://flintlib.org> (2026).
- [70] A. Meurer *et al.*, SymPy: Symbolic computing in Python, *PeerJ Comput. Sci.* **3**, e103 (2017).
- [71] P. Virtanen *et al.* (SciPy 1.0 Contributors), SciPy 1.0: Fundamental algorithms for scientific computing in Python, *Nat. Methods* **17**, 261 (2020).
- [72] J. D. Hunter, Matplotlib: A 2D graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).
- [73] H. Dette and W. J. Studden, Matrix measures, moment spaces and Favard’s theorem for the interval $[0, 1]$ and $[0, \infty)$, *Linear Algebra Appl.* **345**, 169 (2002).
- [74] K. Kobayashi, K. Nakata, and M. Kojima, A conversion of an SDP having free variables into the standard form sdp, *Comput. Optim. Applic.* **36**, 289 (2007).
- [75] T. Blum, Lattice calculation of the lowest order hadronic contribution to the muon anomalous magnetic moment, *Phys. Rev. Lett.* **91**, 052001 (2003).
- [76] R. Aliberti *et al.*, The anomalous magnetic moment of the muon in the Standard Model: An update, *Phys. Rep.* **1143**, 1 (2025).
- [77] S. J. Brodsky and E. De Rafael, Suggested boson-lepton pair couplings and the anomalous magnetic moment of the muon, *Phys. Rev.* **168**, 1620 (1968).
- [78] B. E. Lautrup and E. De Rafael, Calculation of the sixth-order contribution from the fourth-order vacuum polarization to the difference of the anomalous magnetic moments of muon and electron, *Phys. Rev.* **174**, 1835 (1968).
- [79] M. Gourdin and E. De Rafael, Hadronic contributions to the muon g-factor, *Nucl. Phys.* **B10**, 667 (1969).
- [80] C. Bouchiat and L. Michel, La résonance dans la diffusion méson π —méson π et le moment magnétique anormal du méson μ , *J. Phys. Radium* **22**, 121 (1961).