

Semidefinite programs at finite fermion density

Scott Lawrence^{✉*}

Department of Physics, University of Colorado, Boulder, Colorado 80309, USA



(Received 18 January 2023; accepted 10 May 2023; published 23 May 2023)

Semidefinite programs can be constructed to provide a nonperturbative view of the zero-temperature behavior of quantum systems. Past work has found that small semidefinite programs (relative to the dimension of Hilbert space) yield quantitatively accurate estimates for the ground-state energy and expectation values. This paper examines the properties of these semidefinite programs when applied to lattice-regulated field theories exhibiting fermion sign problems. Specifically on the finite-density Thirring model, small semidefinite programs still yield quantitatively accurate results, and there is no indication that these methods encounter exponential costs (analogous to the fermion sign problem) at finite chemical potential.

DOI: [10.1103/PhysRevD.107.094511](https://doi.org/10.1103/PhysRevD.107.094511)

I. INTRODUCTION

Monte Carlo methods, in the framework of lattice field theory, are often the only technique available for non-perturbative computations of expectation values in quantum field theories. These methods generally proceed by transforming a quantum expectation value to an equivalent classical expectation value in one higher space-time dimension. Unfortunately, for a wide variety of systems (notably including relativistic fermions at finite density and the Hubbard model away from half filling), the resulting classical expression possesses a Boltzmann factor which is generally complex, and therefore cannot be viewed as a probability distribution for Monte Carlo sampling. Under standard complexity-theoretic assumptions, a theorem of Troyer and Wiese states that there is no fully general solution to these *fermion sign problems* [1].

Despite the nonexistence of a general solution to fermion sign problems, individual systems may be analyzed with a variety of approaches: complex Langevin [2], the density of states method [3], canonical methods [4,5], reweighting methods [6], series expansions in the chemical potential [7], fermion bags [8], analytic continuation from imaginary chemical potentials [9], and contour deformation methods [10]. These methods are frequently successful for low-dimensional (typically 1 + 1-dimensional) systems, but none have yet been able to provide reliable information about e.g. cold, dense quark matter.

A separate family of approaches to the zero-temperature behavior of quantum systems has recently been developed: the *quantum-mechanical bootstrap* [11–14]. Here we use the observation that, for any quantum mechanical operator \mathcal{O} , the expectation value $\text{Tr}\rho\mathcal{O}^\dagger\mathcal{O}$ must be non-negative in any state ρ . As a direct consequence, if we select a basis of operators $\{\mathcal{O}_i\}$ and a particular state ρ , we may construct a matrix guaranteed to be positive semidefinite¹:

$$M_{ij} \equiv \langle \mathcal{O}_i^\dagger \mathcal{O}_j \rangle \geq 0. \quad (1)$$

That is, for any vector v , we have $v^T M v \geq 0$. When the chosen basis of operators is complete—when any operator in the Hilbert space can be obtained as a linear combination of the elements of the basis—any set of expectation values consistent with Eq. (1) and the system’s commutation relations is obtainable by some quantum state. In other words, every possible inequality describing the quantum system is captured by Eq. (1).

This bound by itself is not usually sufficient to provide interesting information about the ground state: after all, high-temperature states (and indeed far-from-equilibrium states) necessarily obey this bound as well. To do better we must include information about the Hamiltonian of interest. For a chosen Hamiltonian H , the ground state $|\Omega\rangle$ is by definition the state that minimizes $\langle \Omega | H | \Omega \rangle \equiv E_0$. Because the above bound captures every possible inequality, the relation $\text{Tr}\rho H \geq E_0$ can be derived from that bound.

The computational strategy of the quantum-mechanical bootstrap is to use Eq. (1) to derive lower bounds on the expectation value $\langle H \rangle$, which must be obeyed by all states.

¹Here and throughout, the notation $M \geq 0$ is used to denote that a matrix is positive semidefinite.

*scott.lawrence-1@colorado.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.

These are bounds on the ground-state energy, and as more operators are used in the construction of M , the bounds converge to the true ground-state energy. Concretely, we can cast the task of finding the ground-state energy as an optimization problem over the space of expectation values, as follows. For a given basis of operators \mathcal{O}_i , we consider all expectation values of the form $\langle \mathcal{O}_i^\dagger \mathcal{O}_j \rangle$. This defines the positive semidefinite matrix M_{ij} of Eq. (1). The commutation relations furnish a set of linear constraints on the matrix elements of M —that is, a sequence of matrices $C^{(k)}$ such that we require

$$0 = \text{Tr} C^{(k)} M. \quad (2)$$

We also impose that the expectation value of the identity operator is 1. This is equivalent to enforcing the correct normalization of the quantum states being looked at. Note that all of these constraints are linear in the expectation values, which is the space in which our optimization problem takes place.

Finally, as long as the Hamiltonian is obtainable by a linear combination of our chosen basis, there is a matrix G such that $\text{Tr} GM$ represents the expectation value of the Hamiltonian. We seek to minimize this linear combination of elements of M subject to the above constraints. Because this optimization problem is convex, the minimum obtained is guaranteed to be the global minimum, and therefore a lower bound on $\langle H \rangle$ that is valid for all states.

This type of optimization is termed a semidefinite program (SDP). This method provides a rigorous lower bound on the ground-state energy. As a result, this method is essentially dual to variational methods, which provide similarly rigorous upper bounds on the ground-state energy. The combination is particularly powerful when applied to systems for which the ground-state energy is a quantity of central physical concern—for example in the study of quantum-mechanical bound states.

Separate from the explicit bound on the ground-state energy, we also obtain estimates of all expectation values used in the construction of the SDP. In situations where the ground-state energy is not of direct physical interest (most relevantly here, lattice-regularized field theories), these estimates are central to the utility of SDP-based methods.

When the basis of operators is complete, performing this minimization requires manipulating objects of size exponential in the volume of the physical system being studied. The same is true of attempting to exactly diagonalize the Hamiltonian, and so SDP-based methods appear to have no advantage. However, in practice it has been found that operator bases that are far from complete nevertheless result in remarkably tight bounds on simple quantum systems. This was reported early on for few-matrix mechanics [11] and various one-body systems [12]. The power of the quantum-mechanical bootstrap derives from this observation: it is often possible to construct a polynomial-sized

truncated basis for which the bootstrapped bound and many inferred expectation values are close to their true values.

Semidefinite programs have played an important role in quantum field theory before the advent of the quantum-mechanical bootstrap being used in this work. Most prominently, the conformal bootstrap provides the most precise constraints available so far on the critical exponents of many conformal field theories [15–17] (see [18] for a review of these methods). Bootstrap methods applied to the S matrix, historically important [19], have also been recently revived [20,21].

This work applies SDP-based methods to quantum field theories by first regularizing the field theories on a spatial lattice (in the Hamiltonian formalism). The resulting system is a (many-body) quantum-mechanical system to which the methods of [11–13] are directly applicable. A similar approach was first used in [22] to treat the Hubbard model, and demonstrated directly on field theories in [23]. The hierarchy of positive-definiteness constraints so central to all quantum mechanical bootstrap methods is known as the Navascues-Pironio-Acin hierarchy [24,25].

The remainder of the paper is organized as follows. Section II introduces the Thirring model, a frequent test bed of methods targeting the fermion sign problem. The algorithms used to construct semidefinite programs, and extract from them expectation values, are detailed in Sec. III. Because the SDPs related to a space of expectation values, they can be constructed either at finite volumes or directly at infinite volume. The convergence behavior at finite volume, where direct diagonalization of the Hamiltonian is possible for comparison purposes, is examined in Sec. IV. Results at infinite volume are presented in Sec. V. Finally, Sec. VI discusses various aspects of the performance of this method, lays out open questions, and suggests avenues for future work.

II. THIRRING MODEL

The Thirring model is one of many afflicted by the fermion sign problem. Variants of this model have long been used as testbeds for methods to tackle sign problems [2,26–31]. This work is in keeping with that tradition.

A lattice Thirring model with staggered fermions is defined by the Hamiltonian

$$\begin{aligned} H_{\text{Thirring}} = & \sum_x [(-1)^x m c^\dagger(x) c(x) + \mu c^\dagger(x) c(x)] \\ & + \sum_{\langle xy \rangle} \left[(-1)^x \frac{c^\dagger(x) c(y) + c^\dagger(y) c(x)}{2} \right. \\ & \left. - g^2 c(x) c^\dagger(x) c^\dagger(y) c(y) \right]. \end{aligned} \quad (3)$$

Here the second sum is taken over all pairs of adjacent sites on an L -site lattice with periodic boundary conditions.




FIG. 1. Ratio of the fermion mass to the boson mass across couplings. The fermion mass is tuned to $m_F = 0.2$ in all cases, on a ten-site lattice. The dotted gray line marks the point at which $m_B = m_F$, and heuristically, the crossing between the strong and weak coupling regimes.

The operators $c(x)$ are defined to obey anticommutation relations according to

$$\{c(x), c(y)\} = 0, \quad \text{and} \quad \{c^\dagger(x), c(y)\} = \delta_{x,y}. \quad (4)$$

The lattice parameters m , μ , and g^2 are the mass, chemical potential, and coupling respectively. In the continuum limit, this model has emergent Lorentz invariance, and describes a single two-component fermionic field.

For positive g^2 , this model has a repulsive interaction between fermions (attractive between a fermion and anti-fermion), producing a fermion-antifermion bound state analogous to mesons in quantum chromodynamics. The mass ratio $\frac{m_B}{m_F}$ of the boson to the meson may be taken as a measure of the strength of the coupling—at weak coupling $m_B \sim 2m_F$, and studies of the strong-coupling sign problem typically aim for $m_B \sim m_F$ [32]. This ratio $\frac{m_B}{m_F}$ is shown in Fig. 1 across a range of couplings, with the fermion mass always tuned to be $m_F = 0.2$ in lattice units, and the calculation performed by exact diagonalization of the Hamiltonian on a ten-site lattice. Calculations in this paper will typically be performed at $g^2 = 0.5$, well into the strong-coupling regime.

III. ALGORITHM

The basic algorithm to study a quantum system via a semidefinite program proceeds in four steps. First, a basis of operators $\{\mathcal{O}_i\}$ is chosen. We will study the quantum system only by looking at expectation values of the form $\langle \mathcal{O}_i^\dagger \mathcal{O}_j \rangle$; all other properties are to be ignored.

The axioms of our quantum system imply two sets of facts about these expectation values. First, the positive definiteness of the inner product on the Hilbert space implies that the matrix $M_{ij} \equiv \langle \mathcal{O}_i^\dagger \mathcal{O}_j \rangle$ is positive semi-definite. Second, the commutations relations imply certain linear relations between the expectation values.

For example, for any sites x, y , we have $\langle c^\dagger(x)c(y) \rangle = \delta_{xy} - \langle c(y)c^\dagger(x) \rangle$. In the second step of the algorithm, these two sets of constraints are combined to construct a semidefinite program. The quantity to be minimized is the expectation value of the Hamiltonian, which can be expressed as a linear combination of matrix elements of M . The constraints are the linear relations between elements of M derived from commutation relations, and the nonlinear requirement $M \succeq 0$.

Once the SDP has been constructed, the third step of the algorithm is to perform the constrained minimization. For this work, the minimization is done using the MOSEK toolkit [33], which implements an interior-point method to perform the optimization [34].

Finally, in the fourth step we wish to extract physical information from the solution of the SDP. For a simple quantum-mechanical system, the lower bound on the ground state is of intrinsic interest. In the context of (lattice-regularized) field theories, the ground-state energy density at a single value of μ has no physically meaningful continuum limit. For zero-temperature, finite-chemical-potential models, the observable of greatest interest is typically the number density $n(\mu)$. Alternatively, the difference between ground-state energies² at two different values of μ provides the integral of $n(\mu)$ according to

$$E(\mu_i) - E(\mu_f) = \int_{\mu_f}^{\mu_i} n(\mu) d\mu. \quad (5)$$

When solving an SDP, typical algorithms (including the interior-point methods used by MOSEK) return not only the optimal value of the objective function (in this case corresponding to the ground-state energy), but also a value of the positive semidefinite matrix M that obtains that optimum. Any observable encoded in the matrix M can therefore be directly estimated once the optimization is complete. In the case of the Thirring model, because a term $\mu \hat{N}$ appears in the Hamiltonian, the fermion number will always be accessible from the expectation values present in M .

The fermion number could also be obtained, in principle, from solving two SDPs for slightly different values of the chemical potential μ , and computing $\langle \hat{N} \rangle = -\frac{\partial}{\partial \mu} \langle \hat{H} - \mu \hat{N} \rangle$. In fact this yields the same result as simply reading off the value of $\langle \hat{N} \rangle$ from the matrix M in the solution to the SDP. To see this, consider the positive semidefinite matrix M obtained after solution. Generically, this matrix will have exactly one vanishing eigenvalue. The corresponding eigenvector v defines some operator $\hat{B} \equiv v_i \mathcal{O}_i$ such that the expectation value $\langle \hat{B}^\dagger \hat{B} \rangle$ is estimated to be 0. Using the linear part of the SDP (i.e. the commutation relations), it

²Note that in a relativistic theory, care must be taken to define $E(\mu)$ to exclude the contribution from the Fermi sea.




FIG. 2. SDP estimation of ground-state energy and fermion number for the Thirring model on ten sites, with $m = 0.05$ and $g^2 = 0.5$, across the full range of chemical potentials. The left panel shows the estimate of the ground-state energy; the right compares the estimates of fermion number. The four bases used are described in the text.

may be shown that $\hat{B}^\dagger \hat{B} = \hat{H} - E$, where E is the SDP-obtained estimate of the ground-state energy. Now let us consider what happens after perturbing the Hamiltonian by a term $\epsilon \hat{N}$. Again generically, \hat{B} is also perturbed by some term of the same order $\hat{B} \rightarrow \hat{B} + \epsilon \hat{C}$. We now see that (expanding to first order in ϵ)

$$\hat{B}^\dagger \hat{B} + \epsilon(\hat{B}^\dagger \hat{C} + \hat{C}^\dagger \hat{B}) = \hat{H} + \epsilon \hat{N} - E, \quad (6)$$

where again the equality $\hat{B}^\dagger \hat{C} + \hat{C}^\dagger \hat{B} = \hat{N}$ follows from the commutation relations. As a result, the change in the estimated ground-state energy is equal to the expectation value of this operator, which is constrained by the commutation relations to be equal to the estimate $\langle \hat{N} \rangle$.

Typical computational methods (e.g. lattice Monte Carlo methods) arrive at the infinite-volume limit by performing calculations at a sequence of finite volumes and then extrapolating to obtain the infinite-volume result. As discussed in [23], because SDP-based methods work directly with operators and make no reference to the Hilbert space itself, it is natural to perform calculations directly in the infinite-volume limit. This is accomplished with two modifications to the algorithm described above. First, instead of minimizing the Hamiltonian expectation value (which is no longer well-defined), we minimize the Hamiltonian density. Second, we impose translational invariance³ by including in the formulation of the SDP additional linear constraints of the form $\langle \mathcal{O} \rangle = \langle T^\dagger \mathcal{O} T \rangle$ for any unitary translation operator T .

The core of the quantum-mechanical bootstrap is the choice of an incomplete basis of operators—in many ways this parallels the need for a choice of basis when explicitly diagonalizing a truncated Hamiltonian. The time required to solve an SDP scales roughly as the cube of the size of the

³Note that this embeds an additional assumption, namely that translational symmetry is not spontaneously broken in the finite-density ground state of the Thirring model.

basis, so it is critical to keep the basis small. At the same time, a poor choice of basis (including an insufficiently large basis) will not yield a tight bound on the ground-state energy. Finally note that, although the approximation to E_0 obtained via a hierarchy of SDPs is systematically improvable in the sense that, with a sufficiently large basis, the approximated E_0 can be brought within any desired tolerance of the exact result, in principle the cost may be exponential in the precision requested. An efficient algorithm depends on a well-chosen basis.

For this paper, we define a hierarchy of four bases. Each basis is both translationally invariant and Hermitian (meaning that for any operator \mathcal{O} in the basis, \mathcal{O}^\dagger is also present). The first, labeled H0, consists of $c(x)$, $c^\dagger(x)$, and $c^\dagger(x)c(x)$ for all lattice sites x . On an L -site lattice, counting the identity, this has $3L + 1$ elements. Note that although the Hamiltonian is not in the span of this basis, it is in the span of operators obtained by taking a product $\mathcal{O}_i^\dagger \mathcal{O}_j$ of two elements of this basis, and so the SDP is able to provide a nontrivial constraint on the ground-state energy.

The next basis, H1, adds to H0 the L operators of the form $c(x)c^\dagger(x+1)$, and their L complex conjugates. This basis has a total of $5L + 1$ elements.

To the elements of H1 we can add the L operators found in the interaction term of the Hamiltonian: for each site x , $c^\dagger(x)c(x)c(x+1)c^\dagger(x+1)$. The basis obtained in this way is labeled H2.

These first three bases are of course inspired by terms found in the Hamiltonian of Eq. (3). The final basis, termed C1, adds to H2 the L operators of the form $c^\dagger(x)c(x+1)c^\dagger(x+2)c(x+2)$, and their complex conjugates. The size of this largest basis is $8L + 1$.

To check the correctness of these methods, Fig. 2 provides a comparison on a ten-site lattice with ground-state energy information obtained by exact diagonalization of the Hamiltonian. The left panel shows the ground-state energy bound proven via SDP. Although the bound itself is not meaningful, an estimate of the integral of number density

between two chemical potentials may be obtained from the difference of energies at those values of μ . (The extra term proportional to μ ensures that the Fermi sea is not counted in the estimated fermion number.) Alternatively, the right panel compares the estimated fermion number to the exact result. For all four bases, the energy bounds are seen to lie below (or on) the exact value, as they should.

Moreover, significant convergence toward the true value is seen as the basis is expanded. On this lattice, the largest basis (C1) has 81 elements. By comparison, a complete basis of operators on the 1024 states would require 2^{20} operators.

The next section will consider the convergence properties of the SDPs obtained from these bases more carefully. Already from Fig. 2, we can see that the basis H0 provides a qualitatively accurate equation of state everywhere away from the step at $\mu \sim 1.8$. The mechanism for this is visible in the left panel of that figure: for $\mu \in [0, 1.5)$, the error in the energy is roughly constant.

IV. CONVERGENCE

It has previously been demonstrated that semidefinite programs can yield efficient and accurate approximations for the ground-state properties of quantum lattices [22,23]. Other methods, notably lattice Monte Carlo methods, suffer at finite chemical potentials and strong couplings due to the fermion sign problem. The purpose of this section is to investigate to what extent the same is true of the SDPs defined in the previous section.

Figure 3 shows the error in the ground-state energy estimate for all four bases described above. The most striking feature of this figure is that all SDPs perform best at *large* chemical potentials—this is the opposite of the “traditional” behavior of Monte Carlo methods. This success is due to the fact that at large chemical potential, the true ground state of the system is saturated, with all fermionic modes being occupied. Any SDP will, at

sufficiently large μ , prefer this state to all others and therefore agree with the exact result.

As more operators are added to the basis, the SDP approximation improves. Again, the largest and easiest improvements are achieved at larger chemical potentials. In fact, of the bases considered, only the largest (C1) is substantially different at $\mu = 0$.

The news is not all good. Early attempts to resolve the fermion sign problem of lattice QCD and related models frequently struggled to replicate the so-called *silver blaze phenomenon* [35]. (Contour deformations methods—such as those based on Lefschetz thimbles—were a notable exception, e.g. [36].) In the infinite-volume limit, this term refers to the fact that the number density remains constant for some interval of chemical potentials $\mu \in [0, \mu_c)$ before suddenly beginning to increase. In QCD, this increase begins slightly below the proton mass—faulty algorithms for alleviating the sign problem would characteristically begin to increase far too early. We see that the SDPs investigated here mostly (H2 is a striking exception) have a similar flaw: the density begins to increase near the *bare* fermion mass of 0.05, rather than the correct physical mass $m_F \approx 0.20$.

Moreover, at a finite volume, all increases in the fermion number ought to be discontinuous jumps: the fermion number in a nondegenerate eigenstate of the Hamiltonian is always an integer. The semidefinite programs do not generally have this behavior, instead appearing to smooth out these jumps. Again, the basis H2 is a striking exception, indicating that there may be some condition on the basis which is sufficient to enforce this desired behavior.

The performance of the SDPs as coupling is changed is more consistent. The left panel of Fig. 3 shows the error in the estimate of the ground-state energy at $\mu = 0.5$, for a sequence of theories of increasing coupling g^2 . All computations are done on a ten-site lattice, with the bare mass m tuned such that the renormalized fermion mass (measured on the same lattice) is $m_F = 0.2$. As the coupling is increased, the SDP’s error increases roughly linearly.




FIG. 3. Convergence of SDPs constructed from the four bases described in the text, across chemical potentials and couplings. The left panel shows the error in the ground-state energy estimate as g^2 is changed, on ten-site lattices, with the bare mass tuned to be $m_F = 0.2$, and at $\mu = 0.5$. All calculations in the right panel are performed at $g^2 = 0.5$ on a ten-site lattice, with a bare mass of $m = 0.05$.

All bases considered give an exact bound for the free theory. For this theory, the raising and lowering operators that diagonalize the Hamiltonian can be constructed via the discrete Fourier transform of the position-space creation and annihilation operators $c^\dagger(x), c(x)$. Therefore, these operators diagonalizing the Hamiltonian are in the span even of the smallest basis, H_0 . This is sufficient for the SDP estimate to be exact.

V. INFINITE VOLUME

As noted in [22], because the semidefinite programs constraining expectation values work in terms of operators rather than configurations or states in Hilbert space, these calculations can be performed directly in the infinite volume limit. Of course, even for a system (like the Thirring model) with a finite-dimensional local Hilbert space, the set of operators available in this limit is infinite. As a result, any finite basis chosen to construct the SDP will yield only an approximation. Nevertheless, each estimate of the ground-state energy density will be a rigorous lower bound to the infinite-volume result, and we may hope that the estimates of the number density provide a good approximation in practice.

In the infinite-volume limit, it is no longer possible to speak of the expectation value of the Hamiltonian. Instead, we must work with a Hamiltonian density $\hat{h}(0)$, defined so that the Hamiltonian is $\hat{H} = \sum_r \hat{h}(2r)$:

$$\begin{aligned} \hat{h}(x) = & m(c^\dagger(x)c(x) - c^\dagger(x+1)c(x+1)) + \frac{1}{2}(c^\dagger(x)c(x+1) + \text{H.c.}) - \frac{1}{2}(c^\dagger(x+1)c(x+2) + \text{H.c.}) \\ & - g^2(c(x)c^\dagger(x)c^\dagger(x+1)c(x+1)c(x+1)c^\dagger(x+1)c^\dagger(x+2)c(x+2)) - \mu(c^\dagger(x)c(x) + c^\dagger(x+1)c(x+1)). \end{aligned} \quad (7)$$

The optimization objective for the SDP will be the expectation value $\langle \hat{h}(0) \rangle$. In order to close out the SDP, we must include some information about how the energy density around site $x = 0$ is related to the energy density elsewhere on the lattice. In this case we will assume translational invariance. In particular, for any operator \mathcal{O} and even displacement $2r$ ($r \in \mathbb{Z}$), we assume that

$$\langle \mathcal{O}(x) \rangle = \langle \mathcal{O}(x + 2r) \rangle. \quad (8)$$

All linear constraints of this form are added to the SDP.

At finite volume, the free theory is exactly solved by an SDP constructed from a basis consisting only of all available creation and annihilation operators. At infinite volume that basis already must be truncated in order to obtain a finite SDP. Any such truncation will have some maximum radius of operators L , and will be an approximation that improves as L is taken to be larger.

The difference between the truncated SDP and the exact result can be thought of as a finite-volume effect. Consider




FIG. 4. Number density estimated via SDPs for free theories. The basis used to construct the SDP, described in the text, grows linearly with the parameter L . For all bare masses the SDP converges to the exact number density, but the convergence is somewhat faster for larger masses, where only short-distance correlations are relevant. All calculations are performed at $\frac{\mu}{m} = \frac{4}{3}$. To make the differences in convergence more apparent, the bottom panel shows the relative error on a log scale.

an SDP constructed from the operators $a(-L), \dots, a(L)$ and their complex conjugates—or, in the case of a theory at nonzero coupling, an SDP consisting of all possible operators with support only on $[-L, L]$. Such an SDP captures all quantum mechanical constraints on expectation values that arise on that range, but has no information about the boundary conditions. As a result, the obtained bound on the energy density can only be as tight as the lowest energy density achievable with any boundary conditions.

The performance of ever-larger SDPs for the free theory is shown in Fig. 4. Each SDP is constructed only from the operators $a(-L), \dots, a(L+1)$, and their complex conjugates. The SDP-estimated number density is compared to the exact lattice result

$$n_{\text{exact}}(\mu) = \frac{2}{\pi} \sin^{-1} \sqrt{\mu^2 - m^2}. \quad (9)$$

The most notable feature is that the approximation converges much faster, as a function of L , when the fermion

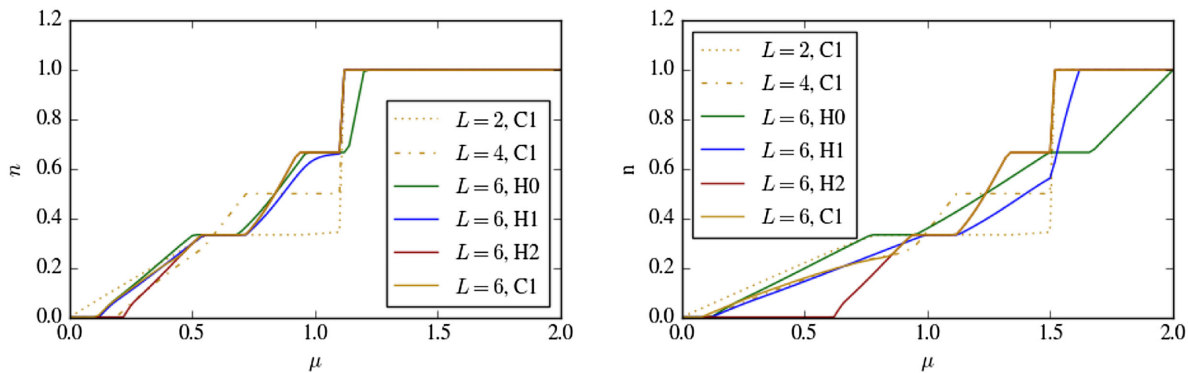


FIG. 5. Number density as a function of chemical potential in the Thirring model, in the infinite volume limit, at weak ($g^2 = 0.1$, left) and strong ($g^2 = 0.5$, right) couplings. Both are simulated at a bare fermion mass $m = 0.05$. Note that the qualitative agreement between different approximations improves as the basis of operators is enlarged.

mass is larger. Physically, this is a result of the fact that finite volume effects decay parametrically with the dimensionless mL .

To perform calculations with a nonzero interaction, it is necessary to return to larger bases of the sort considered in the previous two sections. Infinite-volume equivalents of the bases H0, H1, H2, and C1 are defined by simply requiring $-L \leq x \leq L + 1$. Figure 5 shows calculations performed in the infinite volume limit for interacting theories. The left panel shows a weakly coupled theory ($g^2 = 0.1$), and the right a strongly coupled theory ($g^2 = 0.5$).

In both theories, the number of visible plateaus increases as the permitted separation L of operators is increased. This again shows the manner in which the finite- L effects are analogous to finite-volume effects: calculations at finite volume will also show a number of plateaus proportional to the volume of the system.

The weakly coupled theory in particular shows good qualitative agreement between the various approximations. Although no extrapolation is performed here, both show signs of convergence as larger bases are considered.

VI. DISCUSSION

This paper has explored the performance of SDP-based methods in analyzing the Thirring model in two spacetime dimensions, at finite fermion density. The SDPs used are small relative to the size of the Hilbert space of the system under study, and there is no indication that the method encounters any obstacles analogous to a fermion sign problem. While Monte Carlo methods perform best at $\mu = 0$ and worst when μ is large enough for the lattice to be saturated, the SDPs have the reverse behavior, with tight bounds most easily achieved for chemical potentials at or near lattice saturation density.

Fermion sign problems are most often alleviated while remaining within a Monte Carlo framework. It is difficult to construct a fair performance comparison between SDPs and these Monte Carlo-based methods. For example, if we

consider the asymptotics of the algorithm when the answer is required to be accurate to within ϵ , we see that SDP methods formally outperform Monte Carlo methods. For a system with a finite-dimensional Hilbert space, a Monte Carlo method requires time exponential in the requested number of significant figures. Since, for a fixed finite-dimensional system, there is an SDP that obtains the exact answer, the time required for the SDP to perform a computation to within that precision is of order 1, formally⁴—but that observation provides us with no practical information about the performance of SDPs.

The SDPs considered here provide qualitatively accurate equations of state at finite volumes, across the entire range of chemical potentials. However, the equations of state obtained fail a few important tests. Most notably, the estimated number density, and the “silver blaze” phenomenon at small chemical potentials is not correctly reproduced.

As with other methods along the lines of the “quantum-mechanical bootstrap” [11–14,23], the method in this paper is in a sense dual to the usual variational principle. Whereas a variational study of a quantum-mechanical system yields a rigorous upper bound on the energy, the solution to an SDP yields a rigorous lower bound to the same quantity. As a result, and as demonstrated explicitly in [23], the two may be profitably used in conjunction to yield both rigorous bounds on the energy, and a heuristic estimate of how well either algorithm is doing at approximating the true ground state.

In the course of this work, the main obstacle encountered to using larger SDPs was the memory required to store and manipulate the SDP, rather than the time required during solving. The largest SDP considered was based on an 81×81 positive semidefinite matrix, requiring around 2 GB and 90 seconds⁵ to solve. It has been noted before [23] that improved software might enable larger bases to be used;

⁴Strictly speaking, there is a term logarithmic in ϵ to account for floating-point precision.

⁵Timed on a 4-core AMD Ryzen 3 5300U.

indeed, major advances in the study of the conformal bootstrap over the last decade were made possible by the development of a specialized software package optimized for the sorts of convex optimization problems encountered in that context [16]. The fact that the results in this paper are constrained more by memory than by time—and that the memory use is dominated by bookkeeping for the interior point method, rather than the size of the matrix itself—suggests that similar advances in the quantum mechanical bootstrap might be achieved by making a different memory-time tradeoff.

It is also possible to work with dramatically larger SDPs by obtaining only an approximate solution [37,38]. This sacrifices the rigor of the obtained lower bound. However, for field theory calculations the lower bound on the energy is not physically relevant, potentially making this a worthwhile tradeoff.

SDP-based methods in the infinite volume limit display a phenomenon analogous to finite-volume effects due to the finite spatial extent of operators used in the basis. This suggests that standard techniques for extrapolating to the infinite-volume limit may be helpful in improving the accuracy of the approximation. This is particularly applicable to field theories, where the fact that an SDP yields a

bound on the energy (rather than just an estimate) is not particularly helpful.

This work has indicated that SDPs may represent a practical approach to analyzing the ground-state behavior of fermions at finite density. The application of SDPs to pure-gauge systems has been explored in [39] (although in the path-integral, rather than Hamiltonian, formalism). A natural next step, therefore, is to test these methods on a gauge theory with a finite fermionic chemical potential.

Code and data availability. All code for this paper is openly available at Ref. [40]. SDPs are solved using the MOSEK toolkit [33].

ACKNOWLEDGMENTS

I am grateful to Paul Romatschke for a variety of helpful discussions over the course of this work, and to Paulo Bedaque, Frederic Koehler, Henry Lamm, and Julian Lenz for comments on drafts of this manuscript. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics program under Award No. DE-SC-0017905.

-
- [1] M. Troyer and U.-J. Wiese, *Phys. Rev. Lett.* **94**, 170201 (2005).
 - [2] G. Aarts and I.-O. Stamatescu, *J. High Energy Phys.* **09** (2008) 018.
 - [3] K. Langfeld and B. Lucini, *Acta Phys. Pol. B Proc. Suppl.* **9**, 503 (2016).
 - [4] A. Alexandru, M. Faber, I. Horvath, and K.-F. Liu, *Phys. Rev. D* **72**, 114513 (2005).
 - [5] P. de Forcrand and S. Kratochvila, *Nucl. Phys. B, Proc. Suppl.* **153**, 62 (2006).
 - [6] Z. Fodor and S. D. Katz, *Phys. Lett. B* **534**, 87 (2002).
 - [7] C. R. Allton, S. Ejiri, S. J. Hands, O. Kaczmarek, F. Karsch, E. Laermann, C. Schmidt, and L. Scorzato, *Phys. Rev. D* **66**, 074507 (2002).
 - [8] S. Chandrasekharan, *Eur. Phys. J. A* **49**, 90 (2013).
 - [9] P. de Forcrand and O. Philipsen, *J. High Energy Phys.* **01** (2007) 077.
 - [10] A. Alexandru, G. Basar, P.F. Bedaque, and N.C. Warrington, *Rev. Mod. Phys.* **94**, 015006 (2022).
 - [11] X. Han, S. A. Hartnoll, and J. Kruthoff, *Phys. Rev. Lett.* **125**, 041601 (2020).
 - [12] D. Berenstein and G. Hulsey, [arXiv:2108.08757](https://arxiv.org/abs/2108.08757).
 - [13] D. Berenstein and G. Hulsey, *J. Phys. A* **55**, 275304 (2022).
 - [14] C. O. Nancarrow and Y. Xin, [arXiv:2211.03819](https://arxiv.org/abs/2211.03819).
 - [15] F. Kos, D. Poland, D. Simmons-Duffin, and A. Vichi, *J. High Energy Phys.* **08** (2016) 036.
 - [16] D. Simmons-Duffin, *J. High Energy Phys.* **06** (2015) 174.
 - [17] R. Rattazzi, V. S. Rychkov, E. Tonni, and A. Vichi, *J. High Energy Phys.* **12** (2008) 031.
 - [18] D. Poland, S. Rychkov, and A. Vichi, *Rev. Mod. Phys.* **91**, 015002 (2019).
 - [19] G.F. Chew, *The S-matrix theory of strong interactions* (W. A. Benjamin, Inc., New York, 1981).
 - [20] Y. He, A. Irrgang, and M. Kruczenski, *J. High Energy Phys.* **11** (2018) 093.
 - [21] S. Caron-Huot and V. Van Duong, *J. High Energy Phys.* **05** (2021) 280.
 - [22] T. Barthel and R. Hübener, *Phys. Rev. Lett.* **108**, 200404 (2012).
 - [23] S. Lawrence, [arXiv:2111.13007](https://arxiv.org/abs/2111.13007).
 - [24] M. Navascués, S. Pironio, and A. Acín, *Phys. Rev. Lett.* **98**, 010401 (2007).
 - [25] M. Navascués, S. Pironio, and A. Acín, *New J. Phys.* **10**, 073013 (2008).
 - [26] A. Alexandru, G. Basar, and P. Bedaque, *Phys. Rev. D* **93**, 014504 (2016).
 - [27] A. Alexandru, G. Basar, P.F. Bedaque, G.W. Ridgway, and N.C. Warrington, *J. High Energy Phys.* **05** (2016) 053.
 - [28] A. Alexandru, P.F. Bedaque, H. Lamm, and S. Lawrence, *Phys. Rev. D* **96**, 094505 (2017).

- [29] A. Alexandru, P. F. Bedaque, H. Lamm, and S. Lawrence, *Phys. Rev. D* **97**, 094510 (2018).
- [30] H. Fujii, S. Kamata, and Y. Kikukawa, *J. High Energy Phys.* **12** (2015) 125; **09** (2016) 172(E).
- [31] J. M. Pawłowski and C. Zielinski, *Phys. Rev. D* **87**, 094503 (2013).
- [32] A. Alexandru, G. Basar, P. F. Bedaque, G. W. Ridgway, and N. C. Warrington, *Phys. Rev. D* **95**, 014502 (2017).
- [33] MOSEK-ApS, MOSEK Optimizer API for PYTHON 9.3.10 (2021).
- [34] F. Alizadeh, *SIAM J. Opt.* **5**, 13 (1995).
- [35] T. D. Cohen, *Phys. Rev. Lett.* **91**, 222001 (2003).
- [36] A. Alexandru, G. Başar, P. F. Bedaque, H. Lamm, and S. Lawrence, *Phys. Rev. D* **98**, 034506 (2018).
- [37] S. Burer and R. D. Monteiro, *Math. Program.* **95**, 329 (2003).
- [38] A. Yurtsever, J. A. Tropp, O. Fercoq, M. Udell, and V. Cevher, *SIAM J. Math. Data Sci.* **3**, 171 (2021).
- [39] P. D. Anderson and M. Kruczenski, *Nucl. Phys.* **B921**, 702 (2017).
- [40] <https://gitlab.com/s.lawrence/bootstrap>.