

A SPECTRAL TRANSFORM METHOD FOR SINGULAR STURM–LIOUVILLE PROBLEMS WITH APPLICATIONS TO ENERGY DIFFUSION IN PLASMA PHYSICS*

JON WILKENING[†] AND ANTOINE CERFON[‡]

Abstract. We develop a spectrally accurate numerical method to compute solutions of a model PDE used in plasma physics to describe diffusion in velocity space due to Fokker–Planck collisions. The solution is represented as a discrete and continuous superposition of normalizable and non-normalizable eigenfunctions via the spectral transform associated with a singular Sturm–Liouville operator. We present a new algorithm for computing the spectral density function of the operator that uses Chebyshev polynomials to extrapolate the value of the Titchmarsh–Weyl m -function from the complex upper half-plane to the real axis. The eigenfunctions and density function are rescaled, and a new formula for the limiting value of the m -function is derived to avoid amplification of roundoff errors when the solution is reconstructed. The complexity of the algorithm is also analyzed, showing that the cost of computing the spectral density function at a point grows less rapidly than any fractional inverse power of the desired accuracy. A WKB analysis is used to prove that the spectral density function is real analytic. Using this new algorithm, we highlight key properties of the PDE and its solution that have strong implications on the optimal choice of discretization method in large-scale plasma physics computations.

Key words. Sturm–Liouville theory, continuous spectrum, Titchmarsh–Weyl m -function, spectral density function, Fokker–Planck collisions, WKB approximation

AMS subject classifications. 34B24, 34L16, 35Q84

DOI. 10.1137/130941948

1. Introduction. Partial differential equations involving singular Sturm–Liouville operators with continuous spectra arise frequently in computational physics. Common approaches to solving them include domain truncation, which often regularizes the operator and makes the spectrum discrete, or projection onto finite dimensional orthogonal polynomial or finite element subspaces, which also leads to discrete spectra. Here we develop an alternative approach in which the continuous spectrum is treated analytically via a spectral transform, and the numerical challenge is in accurately representing and evaluating the integrals giving the exact solution.

While the methods developed in this paper to diagonalize singular Sturm–Liouville operators are quite general, we will describe them in the context of velocity-space diffusion in one dimension,

$$(1.1) \quad \frac{\partial h_a}{\partial t} = \frac{1}{v^2} \frac{\partial}{\partial v} \left[\tilde{\Psi}(v) \left(2v^2 h_a + v \frac{\partial h_a}{\partial v} \right) \right] \quad (v > 0, t > 0),$$

where $\tilde{\Psi}(v) = [\operatorname{erf}(v) - v \operatorname{erf}'(v)]/(2v^2)$ is the Chandrasekhar function and $\operatorname{erf}(v) = 2\pi^{-1/2} \int_0^v e^{-r^2} dr$ is the error function. The diffusion operator on the right-hand side

*Received by the editors October 18, 2013; accepted for publication (in revised form) January 6, 2015; published electronically March 19, 2015.

<http://www.siam.org/journals/siap/75-2/94194.html>

[†]Department of Mathematics and Lawrence Berkeley National Laboratory, University of California, Berkeley, CA 94721 (wilkening@berkeley.edu). This author’s work was supported in part by the U.S. Department of Energy, Office of Science, Applied Scientific Computing Research, under award DE-AC02-05CH11231, and by the National Science Foundation under award DMS-0955078.

[‡]Courant Institute of Mathematical Sciences, New York, NY 10012 (cerfon@cims.nyu.edu). This author’s work was supported in part by the U.S. Department of Energy, Office of Science, Fusion Energy Sciences, under award DE-FG02-86ER53223.

of (1.1) plays an important role in numerical simulations of systems governed by the Fokker–Planck kinetic equation [24, 25]

$$(1.2) \quad \frac{\partial f_a}{\partial t} + \mathbf{v} \cdot \nabla f_a + \frac{q_a}{m_a} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} f_a = \sum_b C(f_a, f_b).$$

Here $f_a(\mathbf{x}, \mathbf{v}, t)$ is the distribution function for particles of species a of charge q_a and mass m_a , and $C(f_a, f_b)$ is the Fokker–Planck collision operator [39, 26, 34] describing collisions between species a and other species, including itself. In many problems of interest, the distribution function f_a is close to a Maxwell–Boltzmann distribution, in which case one often writes $f_a = f_{Ma} + h_a$, where f_{Ma} is the Maxwell–Boltzmann distribution and $h_a \ll f_{Ma}$ [9, 3, 28]. The collision operator $C(f_a, f_b)$ may then be linearized about f_{Ma} [26, 2, 29]. When this is done, the operator on the right-hand side of (1.1) is the term in the linearized version of $C(f_a, f_b)$ that represents energy diffusion resulting from the collision of h_a with a Maxwellian background f_{Ma} [26]. In this context, the variable v in (1.1) is the speed coordinate $|\mathbf{v}|$.

Numerical solution of (1.2) is expensive due to the high-dimensional phase space in which the distribution functions evolve [9, 3, 2]. To reduce the computational time without sacrificing accuracy, it is important to develop optimized discretization techniques [8, 29]. For the speed coordinate, nonclassical orthogonal polynomials [41, 29, 20] are emerging as a promising alternative to finite difference methods [9, 4]. Equation (1.1) is well suited to assess the merits of these discretization schemes for the speed coordinate in (1.2) without the computational overhead of a high-dimensional phase space. It is physically relevant since the right-hand side of (1.1) can be identified as the energy-diffusion term in the “test-particle” part of the linearized Landau collision operator [26, 1], as already discussed. The solution also relaxes to a Maxwellian distribution as $t \rightarrow \infty$, as one would physically expect from a collision operator, and the equation is “mass-conserving”; i.e., the integral $\int_0^\infty 4\pi v^2 h_a(v, t) dv$ remains constant in time. As such, (1.1) and close variants of it are often used as standard tests of the accuracy and conservation properties of numerical solvers, as illustrated by Example IV.A in [2], for instance.

In the present paper, we develop a spectral transform method to study the dynamics of (1.1) in detail. In subsequent work [49], jointly with Landreman, we will study the projected dynamics of this equation in finite-dimensional spaces of orthogonal polynomials. Roughly speaking, we show in this paper how to efficiently evaluate the exact solution by discretizing a continuous transform, while in [49] we discretize the PDE before evolving the solution. The latter approach is faster and better suited to large scale computations of the full Fokker–Planck equation, while the current approach clarifies the role of the continuous spectrum in the dynamics and provides an independent means of validating the orthogonal polynomial approach. Validation is especially important in singular cases where the true solution of (1.1) leaves the finite-dimensional subspace and later returns to a point that may or may not agree closely with the solution of the projected dynamics. These singular cases not only are of academic interest but in fact also correspond to situations of practical interest in plasma physics, such as the calculation of the resistivity of a homogeneous, unmagnetized plasma [29]. The construction of an accurate discretization of the exact solution of (1.1) with singular initial conditions is therefore a useful tool to evaluate the performance and accuracy of numerical solvers for the Fokker–Planck equation, and to understand their behavior. It is also of intrinsic theoretical interest, with applications beyond plasma physics, to be able to diagonalize differential operators with

continuous spectra.

Existing algorithms for computing spectral density functions of singular Sturm–Liouville problems either employ a domain truncation technique (accelerated by Richardson extrapolation) [36, 19, 18] or use formulas for the density function [44, 14, 16, 15] that avoid domain truncation but apply only to Sturm–Liouville problems in standard form, $Lu = -u'' + q(x)u = \lambda u$. Further details about both approaches are given in sections 3 and 5 and in [48]. While it is always possible to reduce to standard form via the Liouville transformation [32], it is often preferable to work with the most physically relevant variables. Our idea is to turn Weyl’s original construction [47] into a numerical algorithm by complexifying the spectral parameter $\lambda = \tau + i\varepsilon$ and extrapolating the Titchmarsh–Weyl m -function to the real axis using Chebyshev interpolating polynomials. In more detail, when λ is complex, $m(\lambda) = -\lim_{x \rightarrow \infty} [u_0(x; \lambda)/u_1(x; \lambda)]$ is the limiting ratio of two solutions of $Lu = \lambda u$. We show that this limit converges exponentially fast relative to the work required to compute the solutions u_0 and u_1 , with a decay rate proportional to $\text{Im}\{\sqrt{\lambda}\}/|\lambda|^{1/2}$ (as a function of work). Making use of arbitrary-order ODE solvers and interpolation methods, the cost of the new algorithm for computing the spectral density function $\rho'(\tau) = (1/\pi) \lim_{\varepsilon \rightarrow 0^+} \text{Im}\{m(\tau + i\varepsilon)\}$ to a tolerance δ grows slower than $\delta^{-\gamma}$ (for any $\gamma > 0$) as $\delta \rightarrow 0$. Proving this requires precise information about the asymptotic behavior of solutions of $Lu = \lambda u$ for large x . We present a WKB analysis in Appendix C and a proof of analyticity of $\rho'(\lambda)$ in Appendix D for the Sturm–Liouville operator associated with (1.1).

In addition to developing a new algorithm for computing spectral density functions of singular Sturm–Liouville problems of the form $Lu = \lambda u$, we show how to use them to evaluate the solution of $u_t = -Lu$ at any later time. Through appropriate changes of variables, the spectral transform of the solution, $\hat{u}(\lambda, t)$, can be represented accurately and concisely as a trigonometric polynomial. The solution $u(x, t)$ in physical space is then expressed as an oscillatory integral in λ . For some initial conditions, the spectral transform $\hat{u}(\lambda, t)$ decays exponentially in λ for $t > 0$ but only algebraically at $t = 0$. Thus, with limited computational resources, the solution of $u_t = -Lu$ often cannot be resolved to the desired level of accuracy until t surpasses a critical value, t^* , where the decay rate of $\hat{u}(\lambda, t^*)$ becomes fast enough.

Remarkably, the same is true of the projected dynamics in some spaces of orthogonal polynomials [49]. For singular initial conditions, the projected dynamics is a poor approximation of the true solution initially, regardless of which space of polynomials is used to represent the solution. However, the true solution will generally return (very nearly) to the space once t exceeds some t^* . For the class of orthogonal polynomials introduced by Shizgal [41] and Landreman and Ernst [29], the true solution agrees with the projected dynamics to 29 digits of accuracy for $t > t^*$ (in quadruple-precision arithmetic). By contrast, for classical Hermite polynomials, it only agrees to 2–3 digits of accuracy for similar computational work. Thus, in one case the projected dynamics evolves to the correct state when t reaches t^* , while in the other case it does not. The methods of the current paper were developed in order to quantify these errors and understand these results.

2. Preliminaries. Our goal is to develop a spectral representation for solutions of the PDE (1.1). For notational convenience, we will use the variable x instead of v for the speed coordinate. The equation may then be written

$$(2.1) \quad \frac{\partial h_a}{\partial t} = \frac{1}{x^2} \frac{\partial}{\partial x} \left[\Psi(x) x^2 e^{-x^2} \frac{\partial}{\partial x} (e^{x^2} h_a) \right] \quad (x > 0, t > 0),$$

where

$$(2.2) \quad \Psi(x) = \frac{1}{2x^3} \left[\operatorname{erf}(x) - \frac{2}{\sqrt{\pi}} x e^{-x^2} \right], \quad \operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-s^2} ds.$$

Several properties of $\Psi(x)$, which differs from the Chandrasekhar function $\tilde{\Psi}(x)$ in the introduction by a factor of x , are established in Lemma B.1 of Appendix B. In particular, $\Psi(0) = 2/(3\sqrt{\pi}) \approx 0.3761$, $\Psi'(0) = 0$, $\Psi(x)$ is monotonically decreasing for $x \geq 0$, and $\Psi(x) \sim (2x^3)^{-1}$ as $x \rightarrow \infty$. Furthermore, it is clear that as the ratio of two odd functions, $\Psi(x)$ is even.

We begin by transforming (2.1) into a self-adjoint system. Let

$$(2.3) \quad u(x, t) = h_a(x, t)e^{x^2}.$$

Then u satisfies

$$(2.4) \quad u_t = -Lu, \quad Lu = -\frac{(\Psi w u)'}{w}, \quad w(x) = x^2 e^{-x^2},$$

where $u_t := \partial u / \partial t$ and the prime symbol stands for the derivative with respect to x . The domain D of L can be characterized precisely (see [10] and section 2.3), but it is difficult to show that L is symmetric on all of D directly. Thus, initially, we will work with the set D_1 of bounded, C^2 functions on $(0, \infty)$ with two bounded derivatives. Such functions extend continuously to $x = 0$ with finite limiting value and slope. Since

$$(2.5) \quad \langle Lu, v \rangle = \langle u, Lv \rangle \quad (u, v \in D_1),$$

where $\langle u, v \rangle = \int_0^\infty u(x) \overline{v(x)} w(x) dx$, we see that the (densely defined) restriction operator $L_1 = L \upharpoonright D_1$ is symmetric on the Hilbert space

$$(2.6) \quad \mathcal{H} = L^2(\mathbb{R}_+; w dx) = \left\{ u : \int_0^\infty |u(x)|^2 w(x) dx < \infty \right\}.$$

L is defined as the graph closure of L_1 , which exists since L_1 is symmetric. A well-known theorem [37] asserts that L is self-adjoint iff $\operatorname{Ran}(L_1 + i)$ and $\operatorname{Ran}(L_1 - i)$ are dense in \mathcal{H} , which may be proved by construction of a Green's function (see section 2.3). We note that D_1 must be small enough that L_1 is symmetric but large enough that D_1 and $\operatorname{Ran}(L_1 \pm i)$ are dense in \mathcal{H} . For singular Sturm–Liouville operators, this boils down to imposing the correct boundary conditions at the endpoints.

2.1. Classification of the endpoints. The operator L in (2.4) is singular at $x = 0$ since $w(0) = 0$ and at $x = \infty$ since the domain is unbounded. We now show that L is of limit circle type at $x = 0$ and limit point type at $x = \infty$ [10, 43, 22]. To classify the endpoints, we study the behavior of solutions of

$$(2.7) \quad -(\Psi w u)' = \lambda w u \quad (\lambda \in \mathbb{C})$$

as $x \rightarrow 0$ and $x \rightarrow \infty$. When $\lambda = 0$, the general solution is

$$(2.8) \quad u(x) = \alpha_1 - \alpha_0 \int_1^x \frac{e^{s^2}}{s^2 \Psi(s)} ds \quad (0 < x < \infty).$$

The integrand may be expanded in a Laurent series about $s = 0$ to obtain

$$(2.9) \quad \frac{e^{s^2}}{s^2\Psi(s)} = \frac{3\sqrt{\pi}}{2s^2} + \frac{12\sqrt{\pi}}{5} + O(s^2).$$

Thus,

$$u(x) = \frac{3\sqrt{\pi}}{2x}\alpha_0 + O(1) \quad (x \ll 1).$$

Since $1/x$ belongs to \mathcal{H} in (2.6), all solutions of (2.7) are square-integrable on $(0, 1)$ with weight function $w(x)$ when $\lambda = 0$. Weyl’s theorem [10, 43, 22] states that this is true for all $\lambda \in \mathbb{C}$ if it is true for one λ . Thus, the limit circle case prevails at $x = 0$.

The situation is different at $x = \infty$. Since $\lim_{s \rightarrow \infty} s^3\Psi(s) = 1/2$, there is an $x_0 \in (0, \infty)$ such that

$$(2.10) \quad \frac{e^{s^2}}{s^2\Psi(s)} \geq se^{s^2} \quad (s \geq x_0).$$

It follows that $u(x)$ in (2.8) with $\alpha_1 = 0$ and $\alpha_0 = -1$ satisfies

$$u(x) = \left(u_0 + \int_{x_0}^x \frac{e^{s^2}}{s^2\Psi(s)} ds \right) \geq \left(u_0 + \int_{x_0}^x se^{s^2} ds \right) = \left(u_0 + \frac{e^{x^2} - e^{x_0^2}}{2} \right)$$

for $x \geq x_0$. The function on the right is not square-integrable on (x_0, ∞) with weight function $w(x)$, so neither is u , and the limit point case prevails at $x = \infty$.

2.2. Rescaled variables. The limit circle case requires a boundary condition. It suffices for this to require that solutions of $Lu = \lambda u$ remain bounded at $x = 0$. However, a linearly independent solution (that blows up at the origin) must also be computed to evaluate the Titchmarsh–Weyl m -function, and both of these solutions grow rapidly as $x \rightarrow \infty$. Thus, it is convenient to rescale u and its derivative to avoid overflow in numerical computations. We define

$$(2.11) \quad y(x) = xe^{-x^2/2}u(x), \quad z(x) = \Psi(x)x^2e^{-x^2/2}u'(x)$$

and note that u belongs to \mathcal{H} iff $y \in L^2(0, \infty)$. In terms of y and z , the ODE (2.7) can be rewritten as

$$(2.12) \quad \frac{d\vec{r}}{dx} = A(x)\vec{r}, \quad A(x) := \frac{1}{x} \begin{pmatrix} 1 - x^2 & \Psi(x)^{-1} \\ -\lambda x^2 & x^2 \end{pmatrix}, \quad \vec{r} = \begin{pmatrix} y \\ z \end{pmatrix},$$

which has a singularity of the first kind [10] at $x = 0$. Formal series solutions of (2.12) are therefore convergent, yielding actual solutions. In the present case, the solution may be expanded in a Taylor series, though other problems may require the use of more general Frobenius series or logarithmic sums to obtain a formal solution [5, 10]. Since

$$\Psi(x)^{-1} = \frac{3\sqrt{\pi}}{2} + \frac{9\sqrt{\pi}}{10}x^2 + \frac{153\sqrt{\pi}}{700}x^4 + \dots$$

is even, we see that $A(x) = x^{-1}[A_0 + A_2x^2 + A_4x^4 + \dots]$ where $A_0 = \begin{pmatrix} 1 & 3\sqrt{\pi}/2 \\ 0 & 0 \end{pmatrix}$. It is therefore natural to construct a fundamental matrix with one column even and the other odd:

$$(2.13) \quad \Phi(x) = \left(\vec{r}_0(x), \vec{r}_1(x) \right) = \begin{pmatrix} y_0(x) & y_1(x) \\ z_0(x) & z_1(x) \end{pmatrix}, \quad \begin{pmatrix} \vec{r}_0 = \vec{c}_0 + \vec{c}_2x^2 + \dots \\ \vec{r}_1 = \vec{c}_1x + \vec{c}_3x^3 + \dots \end{pmatrix}.$$

Matching terms yields

$$\begin{aligned} A_0 \vec{c}_0 &= 0, & (2kI - A_0) \vec{c}_{2k} &= \sum_{j=1}^k A_{2j} \vec{c}_{2k-2j} & (k \geq 1), \\ (I - A_0) \vec{c}_1 &= 0, & ((2k + 1)I - A_0) \vec{c}_{2k+1} &= \sum_{j=1}^k A_{2j} \vec{c}_{2k+1-2j} & (k \geq 1), \end{aligned}$$

where I is the 2×2 identity matrix. Since the eigenvalues of A_0 are 0 and 1, nontrivial vectors \vec{c}_0 and \vec{c}_1 exist, which are defined up to multiplicative factors. Once these factors are chosen, the higher order coefficients \vec{c}_{2k} and \vec{c}_{2k+1} are uniquely determined from the recursion relationships given above. The leading terms are

$$(2.14) \quad \left(\vec{r}_0, \frac{\vec{r}_1}{x} \right) = \begin{pmatrix} \frac{3\sqrt{\pi}}{2} & 1 \\ -1 & 0 \end{pmatrix} - \begin{pmatrix} \frac{9}{40}(14\sqrt{\pi} + 5\pi\lambda) & \frac{1}{4}(2 + \sqrt{\pi\lambda}) \\ \frac{1}{4}(2 + 3\sqrt{\pi\lambda}) & \lambda/3 \end{pmatrix} x^2 + O(x^4).$$

The arbitrary constants were chosen so that

$$(2.15) \quad u_1(0) = 1, \quad W[u_0, u_1] = \Psi w(u_0 u_1' - u_1 u_0') = \frac{1}{x} \det \Phi(x) \equiv 1,$$

where $u_j = w^{-1/2} y_j$, $u_j' = x^{-1} w^{-1/2} \Psi^{-1} z_j$ ($j = 0, 1$), and W is the Wronskian determinant. The general solution $u = \alpha_0 u_0 + \alpha_1 u_1$ reduces to (2.8) when $\lambda = 0$.

Note that $\Phi(x)$ is analytic in a complex neighborhood of $x = 0$; i.e., (2.12) has only an apparent singularity [10] at $x = 0$. The determinant of a fundamental matrix is always zero at an apparent singularity, which is true in our case as $\det \Phi(x) = x$. An alternative first-order system using $u(x)$ and $u'(x)$ as components would yield a fundamental matrix with a pole at $x = 0$. Another alternative in which $z(x)$ is replaced by $\Psi(x)y'(x)$ yields the equation

$$(2.16) \quad -(\Psi y')' + V(x)y = \lambda y, \quad V(x) = (1 - x^2) \frac{\Psi'(x)}{x} + (x^2 - 3)\Psi(x),$$

which is self-adjoint and regular at the origin, alleviating the need for initialization with series solutions. This advantage comes at the cost of $V(x)$ being more expensive to evaluate than $A(x)$ in (2.12) due to the additional $\Psi'(x)/x$ term. We also note that the formulas for $\Psi(x)$ and $\Psi'(x)/x$ in (2.16) are numerically unstable near the origin and have to be computed with a series for small x anyway.

2.3. Green’s function and the Titchmarsh–Weyl m -function. For any $\lambda \in \mathbb{C}$ with $\text{Im}\{\lambda\} \neq 0$, we can construct a Green’s function for $L - \lambda$. We seek an operator $[\mathcal{G}(\lambda)f](x) = \int_0^\infty G(x, \xi; \lambda) f(\xi) w(\xi) d\xi$ that satisfies

$$(2.17) \quad \mathcal{G}(\lambda)(L - \lambda)u = u, \quad (L - \lambda)\mathcal{G}(\lambda)f = f$$

for a wide class of functions u and f , which we characterize in detail below. The key to the construction is to identify the complex number $m(\lambda)$, unique in the limit point case, for which

$$(2.18) \quad \chi(x; \lambda) := u_0(x; \lambda) + m(\lambda)u_1(x; \lambda) \text{ belongs to } \mathcal{H}.$$

Although $u_0(x; \lambda) = x^{-1} e^{x^2/2} y_0(x; \lambda)$ has a simple pole at $x = 0$ for all λ , it is square-integrable on $(0, 1)$ because $w(x) = x^2 e^{-x^2}$. Thus, $m(\lambda)$ is determined by the behavior of u_0 and u_1 as $x \rightarrow \infty$. Indeed, when $\text{Im}\{\lambda\} > 0$, one may show [10] that the set of complex numbers $m_b(\lambda)$ for which $\chi(x) = u_0(x; \lambda) + m_b(\lambda)u_1(x; \lambda)$

satisfies $\cos \beta \chi(b) + \sin \beta p(b) \chi'(b) = 0$ forms a nested family of circles (parametrized by β) in the upper half-plane that converge to $m(\lambda)$ as $b \rightarrow \infty$. The parameter β is independent of b , m , and λ and represents a general self-adjoint boundary condition that could be imposed at the right endpoint if the domain were truncated to the finite interval $[0, b]$. Taking $\beta = 0$, we have

$$(2.19) \quad m(\lambda) = - \lim_{b \rightarrow \infty} \frac{u_0(b; \lambda)}{u_1(b; \lambda)} = - \lim_{b \rightarrow \infty} \frac{y_0(b; \lambda)}{y_1(b; \lambda)}.$$

Note that this also holds for $\text{Im}\{\lambda\} < 0$, the only difference being that the nested family of circles is then in the lower half-plane.

Once $m(\lambda)$ and $\chi(x; \lambda)$ are known, the Green’s function may be written as

$$(2.20) \quad G(x, \xi; \lambda) = \begin{cases} u_1(x; \lambda)\chi(\xi; \lambda), & x < \xi, \\ u_1(\xi; \lambda)\chi(x; \lambda), & x > \xi. \end{cases}$$

Suppressing λ to simplify the notation, we have

$$(2.21) \quad \mathcal{G}f(x) = \chi(x) \int_0^x u_1(\xi)f(\xi)w(\xi) d\xi + u_1(x) \int_x^\infty \chi(\xi)f(\xi)w(\xi) d\xi.$$

If f is continuous on $(0, \infty)$ and belongs to \mathcal{H} , it follows from (2.21) and the Wronskian identity that $u = \mathcal{G}f$ has two continuous derivatives on $(0, \infty)$ and satisfies $(L - \lambda)u = f$, where $L - \lambda$ is applied pointwise as a differential operator. It follows that $(L_1 - \lambda)$ from section 2 has a dense range for $\text{Im}\{\lambda\} \neq 0$; i.e., L_1 is essentially self-adjoint. In more detail, let S_1 be the set of continuous functions f with compact support that satisfy $\int_0^\infty u_1 f w d\xi = 0$. If $f \in S_1$, then $\mathcal{G}(\lambda)f \in D_1$ (defined in section 2) and $(L_1 - \lambda)\mathcal{G}(\lambda)f = f$ (already shown). Moreover, S_1 is dense in \mathcal{H} . In particular, the “orthogonality” condition $\langle u_1, f \rangle = 0$ does not preclude density as $u_1 \notin \mathcal{H}$. In fact, if $f \in \mathcal{H}$ is any function for which $\langle u_1, f \rangle$ is finite, then $\langle u_1, f \rangle$ can be adjusted to zero with a continuous, compactly supported perturbation to f of arbitrarily small \mathcal{H} -norm. We leave the details to the reader.

We now characterize the domain D of L . The formula (2.21) is well defined for $f \in \mathcal{H}$ (still assuming $\text{Im}\{\lambda\} \neq 0$). Its derivative exists almost everywhere, where it equals $\chi'(x) \int_0^x u_1 f w d\xi + u_1'(x) \int_x^\infty \chi f w d\xi$, which is locally absolutely continuous. Thus, $\mathcal{G}(\lambda)f(x)$ is actually differentiable everywhere, and $(L - \lambda)\mathcal{G}(\lambda)f = f$ almost everywhere. We claim that $\|\mathcal{G}(\lambda)f\| \leq \|f\|/|\text{Im } \lambda|$; i.e., $\mathcal{G}(\lambda)$ is bounded. This can be seen as follows. Let $f_n \rightarrow f$ in \mathcal{H} with $f_n \in S_1$. A standard argument [38] using the symmetry of L_1 shows that $\|\mathcal{G}(\lambda)f_n\| \leq \|f_n\|/|\text{Im } \lambda|$. Since $\mathcal{G}(\lambda)f_n$ is a Cauchy sequence, it converges to some $g \in \mathcal{H}$ with $\|g\| \leq \|f\|/|\text{Im } \lambda|$. It also converges pointwise to $\mathcal{G}(\lambda)f(x)$, by (2.21). As a result, $\mathcal{G}(\lambda)f = g$ a.e., as required. In addition to belonging to \mathcal{H} , $\mathcal{G}(\lambda)f(x)$ remains finite as $x \rightarrow 0$. Indeed, the first term on the right-hand side of (2.21) approaches zero since $\chi(x) = O(x^{-1})$, and the integral is bounded by $(\int_0^x |u_1(\xi)|^2 w(\xi) d\xi)^{1/2} \|f\| = O(x^{3/2})$. The second term approaches $u_1(0)\langle f, \chi \rangle$, which is finite. By uniqueness of the solutions of $(L - \lambda)u = f$ with $u \in \mathcal{H}$ and u bounded near the origin, the range D of $\mathcal{G}(\lambda)$ consists precisely of those functions $u \in \mathcal{H}$ with a locally absolutely continuous derivative such that $Lu \in \mathcal{H}$. Moreover, (2.17) holds for all $u \in D$, $f \in \mathcal{H}$, $\lambda \notin \mathbb{R}$. This set D , which is independent of λ , is the domain of L .

2.4. Spectral transform. From the general theory of singular self-adjoint eigenvalue problems [10, 44, 43, 22], we know that $m(\lambda)$ (and hence the Green’s function) is analytic on $\mathbb{C} \setminus \mathbb{R}$, with simple poles at the eigenvalues of L and a branch cut across the continuous spectrum of L . The imaginary part of $m(\lambda)$ is positive for $\text{Im } \lambda > 0$, so $m(\lambda)$ is a Pick–Nevanlinna function [11]. Any such function can be represented in the canonical form [11]

$$(2.22) \quad m(\lambda) = A\lambda + B + \int_{-\infty}^{\infty} \left(\frac{1}{s - \lambda} - \frac{s}{s^2 + 1} \right) d\rho(s) \quad (\lambda \in \mathbb{C} \setminus \mathbb{R}),$$

where $A \geq 0$, B is real, and $\rho(s)$ is a nondecreasing, real-valued function such that $\int d\rho(s)/(1 + s^2) < \infty$. Using the fact from section 2.3 that $m(\lambda) = \lim_{b \rightarrow \infty} m_b(\lambda)$, one may show [6] that $A = 0$. Setting $\lambda = \tau + i\varepsilon$, the imaginary part of (2.22) yields

$$(2.23) \quad \text{Im}\{m(\tau + i\varepsilon)\} = \int_{-\infty}^{\infty} \frac{\varepsilon d\rho(s)}{(s - \tau)^2 + \varepsilon^2}.$$

An expression for ρ as a function of τ can be derived from (2.23). Integrating (2.23) from τ_0 to $\tau_1 > \tau_0$, one finds $\int_{-\infty}^{\infty} [\tan^{-1}(\frac{s-\tau_0}{\varepsilon}) - \tan^{-1}(\frac{s-\tau_1}{\varepsilon})] d\rho(s)$ on the right-hand side of the equation. Taking the limit as $\varepsilon \rightarrow 0$, this integrand approaches π for $\tau_0 < s < \tau_1$ and 0 outside this range. It is also positive and bounded by $12\pi(\tau_1 - \tau_0)^2 / [(2s - (\tau_0 + \tau_1))^2 + 3(\tau_1 - \tau_0)^2]$ for $s \in \mathbb{R}$ and $0 < \varepsilon < \tau_1 - \tau_0$. Thus, by the dominated convergence theorem,

$$(2.24) \quad \rho(\tau_1) - \rho(\tau_0) = \lim_{\varepsilon \rightarrow 0^+} \frac{1}{\pi} \int_{\tau_0}^{\tau_1} \text{Im}\{m(\tau + i\varepsilon)\} d\tau$$

at points of continuity τ_0, τ_1 of ρ . Moreover, there is a 1-1 norm-preserving correspondence between $f \in \mathcal{H}$ and $\hat{f} \in L^2(\mathbb{R}; d\rho)$, where one function may be obtained from the other by the relations [10]

$$(2.25) \quad f(x) = \int_{-\infty}^{\infty} \hat{f}(\lambda) u_1(x; \lambda) d\rho(\lambda), \quad \hat{f}(\lambda) = \int_0^{\infty} f(x) u_1(x; \lambda) w(x) dx.$$

The second integral defines $\hat{f}(\lambda)$, the first gives the inversion formula, and $\|\hat{f}\| = \|f\|$. As with the Fourier transform, if f or \hat{f} belongs to L^2 but not L^1 , then the integrals in (2.25) must be defined through a limiting procedure, e.g., by multiplying f or \hat{f} by the characteristic function supported on $[-k, k]$ and letting $k \rightarrow \infty$. Equivalently, the integrals may be regarded as improper integrals rather than Lebesgue integrals over the whole integration domain; see [10] for precise statements.

Since the transform pair (2.25) preserves the L^2 -norm, the jump discontinuities of ρ are precisely the eigenvalues of L , and the size of the jump at the k th eigenvalue is $\|u_1(\cdot; \lambda_k)\|_{\mathcal{H}}^{-2}$. We note that all eigenvalues are simple since any solution of $Lu = \lambda_k u$ that is not a multiple of $u_1(\cdot; \lambda_k)$ will blow up as $x \rightarrow 0^+$. Since $e^{im(\lambda)}$ is bounded in the upper half-plane, standard arguments [42, 40] show that

$$m(\tau^+) := \lim_{\varepsilon \rightarrow 0^+} m(\tau + i\varepsilon)$$

exists and is finite for a.e. $\tau \in \mathbb{R}$. Decomposing $d\rho = \rho'(\tau) d\tau + d\rho_s$ into absolutely continuous and singular components, it follows from (2.23) that for a.e. $\tau \in \mathbb{R}$,

$$(2.26) \quad \rho'(\tau) = \frac{1}{\pi} \text{Im}\{m(\tau^+)\},$$

and $d\rho_s$ is supported on the set $\{\tau : \lim_{\varepsilon \rightarrow 0^+} |m(\tau + i\varepsilon)| = \infty\}$; see [42, 40]. While examples can be constructed in which $d\rho_s$ has a singular continuous component [35], the usual situation [46, 30] is that $d\rho_s$ has a pure point spectrum and $\rho(s)$ is absolutely continuous between eigenvalues. In that case, if there are no eigenvalues between $\tau - a$ and $\tau + a$, (2.22) may be written as a Cauchy integral:

$$(2.27) \quad \frac{m(\lambda)}{2\pi i} = \varphi(\lambda) + \frac{1}{2\pi i} \int_{\tau-a}^{\tau+a} \frac{\rho'(s)}{s-\lambda} ds.$$

The function $\varphi(\lambda)$, which includes B and the remaining portions of the integral in (2.22), is analytic in the upper and lower half-planes as well as in the disk $|\lambda - \tau| < a$. We will show in Appendix D that for L in (2.4), $\rho'(\lambda)$ is real-analytic (for $\lambda > 0$). Thus, if the radius of convergence of $\rho'(\lambda)$ at $\lambda = \tau$ exceeds a , then the contour from $\tau - a$ to $\tau + a$ along the real axis in (2.27) can be deformed to the semicircle $S_{\tau,a}$ in the lower half-plane, and $m(\lambda)$ has an analytic continuation from the upper half-plane to $S_{\tau,a}$. This leads to rapid convergence of polynomial extrapolation methods from the upper half-plane to the real axis, as we will see in section 3.3.

Note that for $\lambda \in \mathbb{R}$, $\rho'(\lambda)u_1(x; \lambda)u_1(\xi; \lambda) = \frac{1}{2\pi i}[G](x, \xi; \lambda)$, where $[G]$ is the jump in the Green’s function across the real λ axis. Some authors [44, 43] make use of this in deriving the transform pair (2.25). It is also useful to know that $\rho(\lambda) = \lim_{b \rightarrow \infty} \rho_b(\lambda)$ at points of continuity of ρ , where ρ_b is the spectral density function for the eigenvalue problem $Lu = \lambda u$ over the finite interval $(0, b)$ with appropriate boundary conditions. Each ρ_b is a right-continuous step function with arbitrary additive constant chosen (in our case) so that $\rho_b(\lambda) = 0$ for $\lambda < 0$.

2.5. The spectrum of L and behavior of solutions of $Lu = \lambda u$. The operator L in (2.4) has only one eigenvalue, $\lambda_0 = 0$. The corresponding eigenvector is $u_1(x; 0) = 1$, which has norm $\frac{1}{2}\pi^{1/4}$. There are no negative eigenvalues since $\langle Lu, u \rangle \geq 0$. In Appendix C, we show that the asymptotic behavior of the general solution of (2.12) for $\lambda > 0$ and $x \gg 1$ has the form

$$(2.28) \quad y(x; \lambda) = Cx^{3/4} \left(1 + \frac{1}{8x\lambda} + \frac{5}{128x^2\lambda^2} + \frac{15}{1024x^3\lambda^3} \right) \cos \left\{ \frac{\sqrt{2x\lambda}}{\lambda^2} \left[\frac{2}{5}x^2\lambda^2 - \frac{x\lambda}{6} - \frac{1}{16} + \frac{1}{64x\lambda} + \frac{5}{3072x^2\lambda^2} + \left(\frac{7}{20480\lambda^3} - \frac{9\lambda}{160} \right) x^{-3} \right] - \theta \right\} + O(x^{-11/4}),$$

where C and θ are constants determined by the initial conditions. Thus, $y_1(x; \lambda)$ does not belong to $L^2(\mathbb{R}_+; dx)$, $u_1(x; \lambda)$ does not belong to \mathcal{H} , and there are no positive eigenvalues. Moreover, a Green’s function does not exist for $\lambda > 0$ since there is no $m(\lambda)$ for which (2.18) holds; thus, the continuous spectrum includes $(0, \infty)$. It actually equals $(0, \infty)$ since (2.22) shows that the Green’s function is analytic across the real axis in regions where $\rho(\lambda)$ is constant, and $\rho(\lambda) = \lim_{b \rightarrow \infty} \rho_b(\lambda) = 0$ for $\lambda < 0$. Alternatively, when $\lambda < 0$, an $m(\lambda)$ for which (2.18) holds can be constructed explicitly since one of the asymptotic solutions $u_{\pm}(x)$ in (C.2) belongs to \mathcal{H} . Thus, $\mathcal{G}(\lambda)$ exists and $L - \lambda$ has a bounded inverse when $\lambda < 0$.

Plots of $y_1(x; \lambda)$ and the error in the asymptotic approximation (2.28) are given in Figure 1 for $\lambda = 0.03$ and $\lambda = 1$. For small λ , the solution exhibits a rapid growth phase before becoming oscillatory. For larger values of λ , the asymptotic formula (2.28) is accurate even for small values of x . The three error curves correspond to the difference between the exact solution (from solving the ODE) and the asymptotic formula (2.28), truncated at orders $x^{-11/4}$, $x^{-7/4}$, and $x^{-3/4}$, respectively. The amplitude C and phase θ in (2.28) were obtained in two stages. First we computed C_k ,

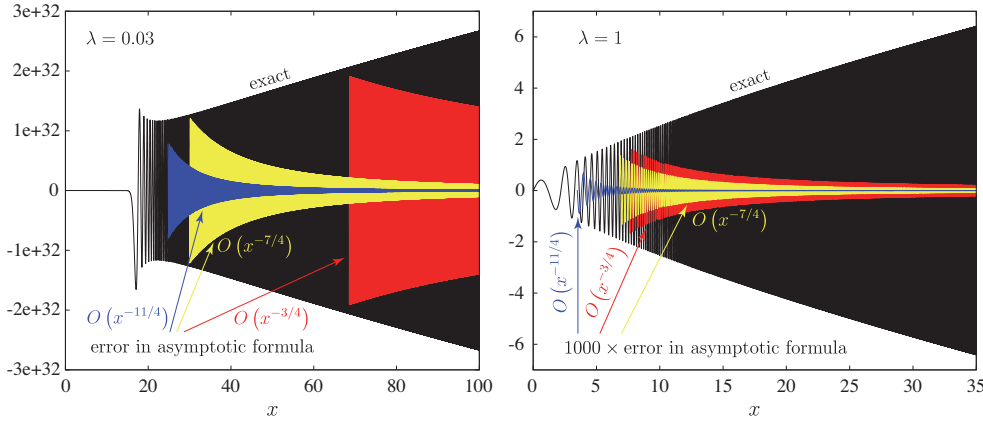


FIG. 1. Solution $y_1(x; \lambda)$ for two values of λ and the error in using the asymptotic formula (2.28) and its lower order variants. The “exact” solution was computed using a 50th order fully implicit Runge–Kutta collocation (IRK) method in quadruple precision arithmetic.

θ_k by fitting the solution through 201 data points near $x_k = 175 + 25k$ for $\lambda = 1$ and $x_k = 500 + 80k$ for $\lambda = 0.03$, with $0 \leq k \leq 5$. Then we extrapolated to $x = \infty$ assuming $C_k \approx C + C_\infty/x_k^4$, $\theta_k \approx \theta + \theta_\infty/x_k^{3.5}$. These values of C and θ were also used for the lower order truncations.

The growth phase observed in Figure 1 occurs only for $0 < \lambda < 0.18704$ and always begins after the first extremum of $y_1(x; \lambda)$. This is because $V(x)$ in (2.16) is negative near the origin, causing $y_1(x; \lambda)$ to execute a small half-oscillation before rapid growth begins (see Figure 2). In more detail, $u_1 = w^{-1/2}y_1$ satisfies

$$(2.29) \quad -\Psi(x)w(x)u_1'(x) = \lambda \int_0^x u_1(s)w(s) ds, \quad u_1(0) = 1.$$

For as long as $u_1(x)$ is positive, the integral is positive and increasing, and $u_1'(x) < 0$ (assuming $\lambda > 0$). Thus, by the mean value theorem, $u_1(x) < 1$ for $0 < x < x_1$, where x_1 is the first zero of $u_1(x)$. This zero exists since we can use (2.29) to bound $u'(x)$ away from zero for $x \in [\varepsilon, x_1]$, where $\varepsilon > 0$ is chosen so $u_1(x) \geq 1/2$ on $[0, \varepsilon]$. As a result, $0 < y_1(x) < xe^{-x^2/2}$ on $(0, x_1)$, so its first extremum is of modest size. The second extremum can be much larger in magnitude if $0 < \lambda < V_{\max}$.

2.6. Spectral representation of the solution. Since the point spectrum of L is $\{0\}$, the singular continuous spectrum is absent (see Appendix D), and the absolutely continuous spectrum is $(0, \infty)$, the transform pair (2.25) simplifies slightly, and the solution of the PDE (2.4) with initial condition $u(x, 0) = f(x)$, $f \in \mathcal{H}$, may be written as

$$(2.30) \quad u(x, t) = \frac{4}{\sqrt{\pi}} \hat{f}(0) + e^{x^2/2} \int_0^\infty [\hat{f}(\lambda)e^{-\lambda t}] \frac{y_1(x; \lambda)}{xY(\lambda)} [Y(\lambda)\rho'(\lambda)] d\lambda,$$

$$(2.31) \quad \hat{f}(\lambda) = \int_0^\infty [xe^{-x^2/2} f(x)] y_1(x; \lambda) dx \quad (\lambda \geq 0),$$

where we used $\hat{u}(\lambda, t) = \hat{f}(\lambda)e^{-\lambda t}$ and $u_1(x; \lambda) = w(x)^{-1/2}y_1(x; \lambda)$ in (2.25). Here $Y(\lambda)$ is an arbitrary scale factor that will be used later to account for the rapid initial growth of $y_1(x; \lambda)$ when λ is small. Evaluation of $\hat{f}(\lambda)$ is straightforward if

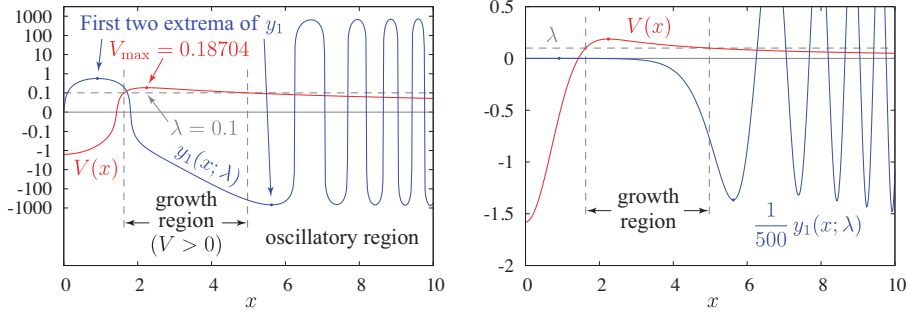


FIG. 2. Plots of $V(x)$ and $y_1(x; \lambda)$, $\lambda = 0.1$, on a signed log (i.e., arcsinh) scale (left) and linear scale (right). Since $\lambda < V_{\max}$, there is a band of x values where $V(x) > \lambda$. In this region, $y_1(x; \lambda)$ grows rapidly as soon as y_1 and y_1' are both negative, since $y_1'' = \Psi^{-1}[(V - \lambda)y_1 - \Psi' y_1']$ remains negative, accelerating the growth rate, until $V - \lambda$ changes sign. Decreasing λ increases the size of the growth region.

$xe^{-x^2/2}f(x)$ decays rapidly (e.g., if f is a polynomial). For a given λ , we just add a third component F to \vec{y} in (2.12) and evolve $dF/dx = xe^{-x^2/2}f(x)y_1(x; \lambda)$ along with y_1 and z_1 until dF/dx is negligible.

3. New algorithm for computing spectral density functions. The most difficult part of computing the solution $u(x, t)$ via (2.30) and (2.31) is determining $\rho'(\lambda)$. A popular approach, implemented in the software package SLEDGE [36, 19, 18], is to compute the step function $\rho_b(\lambda)$ for the regular problem over $(0, b)$ and let $b \rightarrow \infty$. For problems of the form $-u'' + V(x)u = \lambda u$, more sophisticated methods have recently been developed [14, 16, 15] that compute $\rho'(\lambda)$ directly, without computing auxiliary functions ρ_b and taking a limit. Further details are given in section 5 and [48]. See also [13, 17], which focus on solving $-u'' + qu = \lambda u$ when both endpoints are singular. In addition, [13] contains a wealth of information on the history of the problem and provides an insightful review of the literature. There are also many papers on computing eigenvalues and eigenfunctions (rather than continuous spectra and density functions) for singular Sturm–Liouville problems; see, e.g., [23, 31].

3.1. Chebyshev extrapolation and convergence rate. Rather than approximating $\rho(\lambda)$ by $\rho_b(\lambda)$ on a truncated domain or searching for specialized formulas for $\rho'(\lambda)$ that avoid solving ODEs with complex λ , we have developed a simpler approach based directly on (2.19) and (2.26). For a given $\lambda > 0$, we choose a complex line segment Γ of length $\ell > 0$,

$$(3.1) \quad \Gamma = \lambda + i\ell\theta \quad (0 < \theta < 1),$$

and choose n collocation points on Γ consisting of nodes of a Chebyshev–Lobatto quadrature scheme, omitting the node at $\theta = 0$:

$$(3.2) \quad \lambda_k = \lambda + i\ell\theta_k, \quad \theta_k = \frac{1}{2} \left[1 - \cos \left(\frac{k\pi}{n} \right) \right], \quad 1 \leq k \leq n.$$

For each λ_k , we evaluate $m(\lambda_k)$ by computing the limit (2.19). We do this by evolving $\vec{r}_0(x; \lambda_k)$ and $\vec{r}_1(x; \lambda_k)$ simultaneously using an arbitrary (e.g., 50th) order fully implicit Runge–Kutta collocation (IRK) method [21] in double- or quadruple-precision

arithmetic. As shown in Appendix C, there exist two solutions of (2.12) of the form

$$(3.3) \quad y_{\pm}(x; \lambda_k) = x^{3/4} P_0(x; \lambda_k) \exp \left\{ \pm i \frac{2\sqrt{2\lambda_k}}{5} P_1(x; \lambda_k) x^{5/2} \right\} [1 + O(x^{-7/2})]$$

for $x \gg 1$, where $P_0(x; \lambda)$ and $P_1(x; \lambda)$ are defined in (C.3) and approach 1 as $x \rightarrow \infty$. Since $\text{Im} \sqrt{\lambda_k}$ is positive, $y_+(x; \lambda_k)$ decays superexponentially as $x \rightarrow \infty$ while all other linearly independent solutions grow. The function $y_1(x; \lambda_k)$ is guaranteed to grow superexponentially since decaying would cause the corresponding u_1 to be an eigenfunction of L . The same is true of $y_0(x; \lambda_k)$ since there is a related self-adjoint boundary value problem in which u_0 would then be an eigenfunction. The limit (2.19) emerges when the decaying mode in $y_0(x; \lambda_k)$ and $y_1(x; \lambda_k)$ becomes negligible in comparison to the growing mode.

Quantitatively, (3.3) gives an asymptotic estimate for the rate of convergence of $-y_0/y_1$ to m . Expressing $y_j(x; \lambda_k) = c_j^- y_-(x; \lambda_k) + c_j^+ y_+(x; \lambda_k)$ for $j = 0, 1$, we have

$$(3.4) \quad \frac{y_0(x; \lambda_k)}{y_1(x; \lambda_k)} = \frac{c_0^-}{c_1^-} \cdot \frac{1 + (c_0^+ y_+)/ (c_0^- y_-)}{1 + (c_1^+ y_+)/ (c_1^- y_-)} = -m(\lambda_k) \left[1 + o\left(e^{-\text{Im}\{4\sqrt{2\lambda_k}/5\} x^{5/2}} \right) \right]$$

for $x \gg 1$. Here we used the fact that $c_0^- \neq 0$ and $c_1^- \neq 0$, as explained above, as well as the estimate $|(1 + \delta_1)/(1 + \delta_2) - 1| \leq 2|\delta_1| + 2|\delta_2|$ when $|\delta_2| \leq 1/2$, and

$$(3.5) \quad |y_+/y_-| e^{\text{Im}\{4\sqrt{2\lambda_k}/5\} x^{5/2}} = [1 + O(x^{-7/2})] e^{-\text{Im}\{(4/5)\sqrt{2\lambda_k}[P_1(x; \lambda_k) - 1]\} x^{5/2}},$$

which converges to zero as $x \rightarrow \infty$ since $\sqrt{\lambda_k}[P_1(x; \lambda_k) - 1] = \frac{5x^{3/2}}{12} [-\lambda_k^{-1/2} + O(x^{-1})]$, so its imaginary part approaches $+\infty$ as $x \rightarrow \infty$ (since $\text{Im} \lambda_k > 0$). We conclude that the relative error in approximating $m(\lambda_k)$ by $-y_0(x; \lambda_k)/y_1(x; \lambda_k)$ decays extremely rapidly, faster than $\exp(-\text{Im}\{4\sqrt{2\lambda_k}/5\} x^{5/2})$, as $x \rightarrow \infty$.

Since this asymptotic estimate for the convergence rate applies only for large x , it is useful to develop an estimate for the error that can be monitored as the solution is evolved numerically. We find that the number of digits in $y_0(x; \lambda_k)/y_1(x; \lambda_k)$ that remain frozen as x increases is roughly the same as the number of correct digits in the numerically computed Wronskian. Thus, we use the stopping criterion that $|W[u_0, u_1] - 1|$ in (2.15) exceeds 1. At this point, \vec{r}_0 and \vec{r}_1 in (2.13) are linearly dependent to machine precision, and continuing further in the evaluation of the limit (2.19) does more harm than good due to additional roundoff errors. Note that loss of accuracy in the Wronskian does not mean y_0 and y_1 are inaccurate; the catastrophic cancellation of digits occurs when the determinant of Φ is computed.

Once $m(\lambda_k)$ is known for each λ_k in (3.2), we compute the interpolating polynomial $q(\theta)$ satisfying

$$(3.6) \quad q(\theta_k) = m(\lambda_k), \quad 1 \leq k \leq n,$$

and evaluate $\frac{1}{\pi} \text{Im}\{q(0)\}$ to approximate $\rho'(\lambda) = \frac{1}{\pi} \text{Im}\{m(\lambda^+)\}$. The results of this “naive” algorithm, and the improved version described in section 3.2 below, are shown in Figure 3, where we computed $\rho'(\lambda)$ at 768 values of λ of the form

$$(3.7) \quad \lambda_j = e^{\sigma_j}, \quad \sigma_j = -4 + \frac{3j}{128}, \quad 0 \leq j < 768.$$

When λ decreases below 0.2, $\rho'(\lambda)$ begins to decrease rapidly, and typical values of $y_1(x; \lambda)$ grow very large. To account for this, we introduced a scale factor $Y(\lambda)$ in

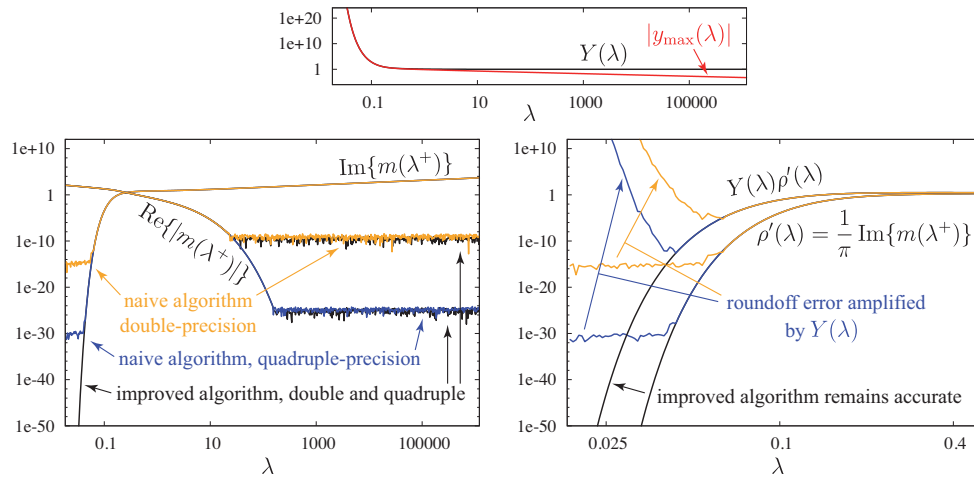


FIG. 3. Plots of the real and imaginary parts of $m(\lambda^+)$, the spectral density function $\rho'(\lambda)$, and its rescaled version, $Y(\lambda)\rho'(\lambda)$, which is used in (2.30) to represent $u(x, t)$. Roundoff errors in the naive algorithm are amplified to unacceptable levels when rescaled.

(2.30), which we define as

$$(3.8) \quad Y(\lambda) = \sqrt{1 + y_{\max}^2(\lambda)},$$

where $y_{\max}(\lambda)$ is the first negative extremum of $y_1(x; \lambda)$, computed using Newton’s method to solve $y_1'(x; \lambda) = 0$ for x , which occurs when $z_1(x; \lambda) = (x^2 - 1)\Psi(x)y_1(x; \lambda)$. For example, in Figure 1, $y_{\max}(0.03) = -1.656 \times 10^{32}$ while $y_{\max}(1) = -0.7377$. The reason we use the first negative extremum (rather than the positive one that precedes it) was explained in Figure 2.

3.2. Avoiding amplification of roundoff error. The poor scaling of $y_1(x; \lambda)$ poses a problem in the reconstruction of the solution via (2.30). Roundoff errors in $\rho'(\lambda)$ near 10^{-15} in double-precision and 10^{-30} in quadruple-precision are amplified to large values by $Y(\lambda)$ when λ is small enough. Fortunately, we are able to improve the algorithm to achieve small relative errors in $\rho'(\lambda)$.

Let us complexify λ again and write $\lambda = \tau + i\varepsilon$ with $\tau > 0$ fixed. We wish to compute $\rho'(\tau)$. First, we write the fundamental matrix in (2.13) as a product,

$$(3.9) \quad \Phi(x; \lambda) = \tilde{\Phi}(x; \lambda)\Phi(x^*; \lambda), \quad \tilde{\Phi}(x^*; \lambda) = I,$$

where x^* is the location of the first negative extremum of $y_1(x; \tau)$, which is independent of ε , and $\tilde{\Phi}(x; \lambda)$ is another fundamental matrix for (2.12), evolved from the identity at $x = x^*$. By analogy with (2.13), we denote the entries of $\tilde{\Phi}$ by $\tilde{y}_0, \tilde{y}_1, \tilde{z}_0, \tilde{z}_1$. When $\varepsilon \neq 0$, the limit (2.19) exists, and we find from (3.9) that

$$(3.10) \quad \begin{aligned} \tilde{m}(\lambda) &:= - \lim_{x \rightarrow \infty} \frac{\tilde{y}_0(x; \lambda)}{\tilde{y}_1(x; \lambda)} = \lim_{x \rightarrow \infty} \frac{y_0(x; \lambda)z_1(x^*; \lambda) - y_1(x; \lambda)z_0(x^*; \lambda)}{y_0(x; \lambda)y_1(x^*; \lambda) - y_1(x; \lambda)y_0(x^*; \lambda)} \\ &= \frac{m(\lambda)z_1(x^*; \lambda) + z_0(x^*; \lambda)}{m(\lambda)y_1(x^*; \lambda) + y_0(x^*; \lambda)}. \end{aligned}$$

Since $\Phi(x^*; \lambda)$ is invertible, the numerator and denominator cannot both be zero, so the limit exists in the extended complex plane. It is in fact finite, for $\tilde{m}(\lambda)$ is the m -function (multiplied by $\exp[(x^*)^2/x^*]$ for the eigenvalue problem $Lu = \lambda u$ on $[x^*, \infty)$

with Dirichlet boundary conditions at the left endpoint. Moreover, $\tilde{m}(\tau^+)$ exists (since $m(\tau^+)$ does) and is given by the final formula of (3.10) with $m(\lambda)$ replaced by $m(\tau^+)$ and λ replaced by τ in y_0, y_1, z_0, z_1 . Next, since $\det \Phi(x^*; \tau) = x^*$, the singular value decomposition of $\Phi(x^*; \tau)$ has the form

$$(3.11) \quad \begin{pmatrix} y_0(x^*; \tau) & y_1(x^*; \tau) \\ z_0(x^*; \tau) & z_1(x^*; \tau) \end{pmatrix} = \begin{pmatrix} a & -b \\ b & a \end{pmatrix} \begin{pmatrix} \alpha\sqrt{x^*} & \\ & \alpha^{-1}\sqrt{x^*} \end{pmatrix} \begin{pmatrix} c & d \\ -d & c \end{pmatrix},$$

where $a^2 + b^2 = 1$, $c^2 + d^2 = 1$, and $\alpha > 0$. We know that if the first singular value is $\alpha\sqrt{x^*}$, the second singular value must be $\alpha^{-1}\sqrt{x^*}$ because the Wronskian is equal to 1. This is convenient because the larger singular value $\alpha\sqrt{x^*}$ can be computed accurately while the smaller one may be severely corrupted by roundoff error. In our algorithm, we ignore the computed version of the smaller singular value and assume it equals $\alpha^{-1}\sqrt{x^*}$ instead. Combining (3.11) with (3.10), we obtain

$$(3.12) \quad \begin{aligned} m &= -\frac{\tilde{m}y_0 - z_0}{\tilde{m}y_1 - z_1} = -\frac{\alpha(b - a\tilde{m})c + \alpha^{-1}(a + b\tilde{m})(-d)}{\alpha(b - a\tilde{m})d + \alpha^{-1}(a + b\tilde{m})(c)} \\ &= -\beta + \frac{1 + \beta^2}{\alpha^2\mu^{-1} + \beta} = \beta^{-1} - \frac{1 + \beta^{-2}}{\alpha^{-2}\mu + \beta^{-1}} \quad \left(\beta = \frac{c}{d}, \quad \mu = \frac{a + b\tilde{m}}{b - a\tilde{m}} \right), \end{aligned}$$

where $m = m(\tau^+)$, $\tilde{m} = \tilde{m}(\tau^+)$, and y_0, z_0, y_1 , and z_1 are evaluated at $(x^*; \tau)$. The resulting formula

$$(3.13) \quad \rho'(\tau) = \begin{cases} \frac{1}{\pi} \operatorname{Im} \left\{ \frac{1 + \beta^2}{\alpha^2\mu^{-1} + \beta} \right\}, & |\beta| \leq 1, \\ -\frac{1}{\pi} \operatorname{Im} \left\{ \frac{1 + \beta^{-2}}{\alpha^{-2}\mu + \beta^{-1}} \right\}, & |\beta| > 1, \end{cases}$$

avoids the cancellation of digits that occurs if α is large and \bar{m} is not simplified in (3.12) to separate out the $-\beta$ (or β^{-1}) term.

The choice of x^* as the first negative extremum of $y_1(x; \tau)$ ensures that $\Phi(x^*; \tau)$ captures the growth phase of $y_1(x; \tau)$ observed in Figures 1 and 2. As a result, $\tilde{\Phi}(x; \lambda)$ begins as the identity matrix in the oscillatory phase of the ODE, where growth is due to λ being complex rather than all solutions growing. This causes \tilde{m} in the improved algorithm to be more accurately computed than m in the naive algorithm. Another advantage of the split (3.9) is that λ can be set to τ when computing $\Phi(x^*; \lambda)$ since it does not depend on x . Of course, λ must remain complex in $\tilde{\Phi}(x; \lambda)$ since the limits $x \rightarrow \infty$ and $\varepsilon \rightarrow 0$ do not commute.

We remark that if $z(x)$ is replaced by $\Psi(x)y'(x)$ and $\Phi(x)$ is evolved according to

$$(3.14) \quad \Phi' = \begin{pmatrix} 0 & 1/\Psi \\ V - \lambda & 0 \end{pmatrix} \Phi,$$

with V as in (2.16), then $\det \Phi(x^*; \tau) = 1$, rather than x^* . Formulas (3.11)–(3.13) remain unchanged, except that $\sqrt{x^*}$ should be omitted from the diagonal in (3.11). This technique of computing $\rho'(\lambda)$ with high relative accuracy works generally and is not tied to L in (2.4). Any convenient choice of x^* can be used as long it is in the “oscillatory” region where the solution grows due to λ being complex. To compute a, b, c, d, α , and \tilde{m} , one may use any first-order system $d\vec{r}/dx = A(x)\vec{r}$ that is equivalent to $Lu = \lambda u$, provided the first component of \vec{r} is a multiple of u . (In our case, $\vec{r} = (y; z)$ with $y = xe^{-x^2/2}u$.) The Abel–Liouville formula $(d/dx) \det \Phi = (\operatorname{tr} A) \det \Phi$ will determine the factor to include with α and α^{-1} in (3.11). Typically, as in (3.14), $\operatorname{tr} A = 0$, and the factor is 1.

3.3. Error bounds, complexity, and optimal parameters. In the improved algorithm, we need to compute $\tilde{m}(\lambda^+)$ for $\lambda > 0$. This is done with an interpolating polynomial $q(\theta)$, just as in (3.6), but matching $\tilde{m}(\lambda_k)$ instead of $m(\lambda_k)$ at θ_k . As in (3.2), λ is now real, while $\lambda_k = \lambda + i\ell\theta_k$, $1 \leq k \leq n$, are complex. We will show in this section that the cost of computing $\tilde{m}(\lambda_k)$ for $k = 1, \dots, n$ is dominated by the $k = 1$ term, regardless of n , and we determine the parameters ℓ and n to best take advantage of the superexponential rate of convergence of $\tilde{y}_0(x; \lambda_k)/\tilde{y}_1(x; \lambda_k)$ to $-\tilde{m}(\lambda_k)$.

3.3.1. Convergence rate. Writing $\tilde{y}_j(x; \lambda_k) = \tilde{c}_j^- y_-(x; \lambda_k) + \tilde{c}_j^+ y_+(x; \lambda_k)$ for $j = 0, 1$, we obtain (3.4) again, but with tildes placed over $y_0, y_1, c_0^\pm, c_1^\pm$, and m . The coefficients \tilde{c}_j^\pm are determined from the initial condition $\tilde{\Phi}(x^*) = I$ by solving

$$(3.15) \quad E(x^*)C = \begin{pmatrix} y_+(x^*) & y_-(x^*) \\ z_+(x^*) & z_-(x^*) \end{pmatrix} \begin{pmatrix} \tilde{c}_0^+ & \tilde{c}_1^+ \\ \tilde{c}_0^- & \tilde{c}_1^- \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

where $z_\pm(x) = \Psi(x)[xy'_\pm(x) + (x^2 - 1)y_\pm(x)]$ and we have dropped λ_k to simplify the notation, since it is fixed in this discussion. We emphasize that $(y_\pm; z_\pm)$ are exact solutions of the ODE (2.12) for which y_\pm have the form (3.3) for $x \gg x^*$. As before, the coefficients \tilde{c}_j^- must be nonzero since there are self-adjoint boundary-value problems, this time on $[x^*, \infty)$, for which \tilde{u}_0 or \tilde{u}_1 would be an eigenfunction with complex eigenvalue if it were a multiple of $u_+ = w^{-1/2}y_+$, the solution that decays as $x \rightarrow \infty$. In other words, neither $y_+(x)$ nor $z_+(x)$ can ever vanish since $\lambda_k \notin \mathbb{R}$. The same conclusions hold if we define $z_\pm(x) = \Psi(x)y'_\pm(x)$ in (3.15) and use (2.16) and (3.14) to evolve the solution instead of (2.12).

The growing solution, $y_-(x)$, is not uniquely determined. Adding any multiple of $y_+(x)$ will not affect the asymptotics of $y_-(x)$ as $x \rightarrow \infty$. Thus, we may assume that the columns of $E(x^*)$ above are orthogonal to each other. Also, $\tilde{\Phi}(x) = E(x)C$ is unchanged if we replace E by ED and C by $D^{-1}C$, where D is diagonal. Thus, we may assume $E(x^*)$ and C are orthogonal matrices in (3.15) if we allow y_+ and y_- to be of the form (3.3) up to constant factors. The signs can be arranged (via D) so that $\tilde{c}_0^+ = -\tilde{c}_1^-$, $\tilde{c}_1^+ = \tilde{c}_0^-$, $y_+(x^*) = -\tilde{c}_1^-$, and $y_-(x^*) = \tilde{c}_0^-$. The argument of (3.4) then gives the relative error estimate

$$(3.16) \quad \frac{1}{|\tilde{m}|} \left| \tilde{m} + \frac{\tilde{y}_0(x)}{\tilde{y}_1(x)} \right| \leq 2\kappa\epsilon(x) = o\left(e^{-\text{Im}\{4\sqrt{2\lambda_k}/5\}[x^{5/2} - (x^*)^{5/2}]}\right),$$

where

$$(3.17) \quad \kappa = 1 + \left| \frac{\tilde{c}_1^-}{\tilde{c}_0^-} \right|^2 = 1 + |\tilde{m}|^{-2}, \quad \epsilon(x) = \left| \frac{y_+(x)/y_+(x^*)}{y_-(x)/y_-(x^*)} \right|.$$

The first inequality in (3.16) is valid once $\epsilon(x) \leq \frac{1}{2}$. Note that $\epsilon(x)$ is the ratio of the magnitudes of the decaying and growing solutions if both are scaled to equal 1 at $x = x^*$, and if the growing solution is chosen to be orthogonal to the decaying solution at $x = x^*$. The $(x^*)^{5/2}$ term in (3.16) accounts for the fact that the error is $O(1)$ when $x = x^*$. It is not strictly necessary but avoids hiding a large constant in the $o(\dots)$ notation that would delay its convergence to zero. As mentioned already, \tilde{c}_0^- and \tilde{c}_1^- are guaranteed to be nonzero, so κ is finite. The results of Figure 4 below show that $\kappa \leq 2$ for $4 \leq \lambda \leq 10^6$ and remains less than 3000 if λ is decreased to 0.01875, assuming $\text{Im } \lambda_k$ is small enough that $|\tilde{m}(\lambda_k)|$ is similar in size to $|\tilde{m}(\lambda^+)|$.

3.3.2. Step count for high order collocation methods. In Appendix A, we show that the number of steps required to evolve (2.12) or (2.16) at λ_k from x_1 to x_2 using a ν -stage Runge–Kutta collocation method of order 2ν while maintaining a bound of δ on the local truncation error satisfies

$$(3.18) \quad N_{\text{steps},k} \leq K_\nu(\delta) \int_{x_1}^{x_2} 5|\lambda_k|^{1/2} \langle x \rangle^{3/2} + 3.5|\lambda_k|^{-1/2} \langle x \rangle^{1/2} dx.$$

Here $K_\nu(\delta) = \max(\frac{3}{2}, \frac{4}{9}\delta^{-1/(2\nu-1)})$, $\langle x \rangle = \sqrt{1+x^2}$, and $\nu \geq 5$ was assumed to obtain these particular constants. For the complexity analysis, it is convenient to decouple x^* from the location where $y_1(x)$ achieves its second extremum. Instead, we set

$$(3.19) \quad x^* = \begin{cases} 1, & \lambda \geq 1, \\ \lambda^{-1}, & 0 < \lambda < 1. \end{cases}$$

This change has little effect on the running time of the algorithm, but is easier to analyze. Recall that we evolve to x^* with λ real to get past the growth phase, then complexify $\lambda_k = \lambda + i\ell\theta_k$, and evolve from x^* until $-y_0/y_1$ converges to $\tilde{m}(\lambda_k)$. As shown in Figure 2, the growth phase occurs in the band of x values for which $V(x) \geq \lambda$. This band exists only when $\lambda < V_{\text{max}} = 0.18704$ and ends well before x reaches $1/\lambda$, since $V(x) < 1/(2x)$ for $x > 0$. Thus, (3.19) is sufficient to traverse the growth phase, when it exists.

Combining (3.18) and (3.19) and making use of $\langle x \rangle^\alpha \leq 1 + x^\alpha$ for $x > 0$, $\alpha \leq 2$, we find that the number of steps required to evolve from x^* to x is bounded by

$$(3.20) \quad \frac{N_{\text{steps},k}}{K_\nu(\delta)|\lambda_k|^{1/2}} \leq \frac{34}{5} [x^{5/2} - (x^*)^{5/2}],$$

where we used $|\lambda_k|^{-1} \leq 1$ when $\lambda = \text{Re } \lambda_k \geq 1$ and $|\lambda_k|^{-1} \leq x$ when $\lambda < 1$ and $x \geq x^* = \lambda^{-1}$. An additional $(77/6) \max(\lambda^{-2}, \lambda^{1/2})K_\nu(\delta)$ steps are required to evolve from 0 to x^* .

3.3.3. Cost of evolving all n solutions from x^* . We see from (3.16) that for small enough δ , the relative error in estimating \tilde{m} by $-\tilde{y}_0/\tilde{y}_1$ will be less than δ if

$$(3.21) \quad b(x) = \ln(1/\delta), \quad \left(b(x) := \text{Im}\{4\sqrt{2\lambda_k}/5\} [x^{5/2} - (x^*)^{5/2}] \right).$$

Here we rely on surpassing the point beyond which $|o(e^{-b(x)})| \leq e^{-b(x)}$ in (3.16). Equivalently, $\epsilon(x)$ and κ in (3.17) must satisfy $\epsilon(x) \leq e^{-b(x)}/(2\kappa)$ by the time $b(x)$ reaches $\ln(1/\delta)$. This happens rapidly since $\epsilon(x)e^{b(x)}$ converges to zero superexponentially in spite of the rapid growth of $e^{b(x)}$, as in (3.5). Combining (3.20) and (3.21), the number of steps required to achieve a relative error of δ satisfies

$$\frac{5}{34} \cdot \frac{4\sqrt{2}}{5} \cdot \frac{N_{\text{steps},k}}{K_\nu(\delta) \ln(1/\delta)} \leq \frac{\sqrt{|\lambda + i\ell\theta_k|}}{\text{Im}\{\sqrt{\lambda + i\ell\theta_k}\}} = 2\frac{\lambda}{\ell\theta_k} + \frac{3}{4}\frac{\ell\theta_k}{\lambda} + O\left(\frac{\ell\theta_k}{\lambda}\right)^3.$$

We see that for fixed δ , $N_{\text{steps},k}$ is a function of $\ell\theta_k/\lambda$, and, to leading order, scales inversely with it. Retaining only this leading term, which is particularly accurate for those θ_k closest to zero (which matter most), we can estimate the total cost (in steps) of computing $\tilde{m}(\lambda_1), \dots, \tilde{m}(\lambda_n)$:

$$(3.22) \quad \frac{\sqrt{2}}{17} \cdot \frac{N_{\text{tot}}}{K_\nu(\delta) \ln(1/\delta)} \lesssim \frac{\lambda}{\ell} \sum_{k=1}^n \theta_k^{-1} = \frac{\lambda}{\ell} \left(\frac{2n^2 + 1}{3} \right).$$

The final equality follows from observing that the polynomial

$$P(\theta) = \theta(\theta_1 - \theta)(\theta_2 - \theta) \cdots (\theta_n - \theta)$$

satisfies $P'(0) = \theta_1\theta_2 \cdots \theta_n$ and $P''(0) = -2P'(0)(\theta_1^{-1} + \cdots + \theta_n^{-1})$. Setting $Q(\vartheta) = P((1 - \cos \vartheta)/2)$, we find (from the choice of the θ_k in (3.2) as Chebyshev–Lobatto nodes) that $Q(\vartheta) = \sin n\vartheta \sin \vartheta$. This gives

$$P'(0) = \lim_{\vartheta \rightarrow 0} \frac{2Q'(\vartheta)}{\sin \vartheta}, \quad P''(0) = \lim_{\vartheta \rightarrow 0} \left(\frac{4Q''(\vartheta)}{\sin^2 \vartheta} - \frac{4Q'(\vartheta)}{\sin^3 \vartheta} \cos \vartheta \right),$$

which evaluate to $4n$ and $-\frac{8}{3}n(2n^2 + 1)$, respectively. The result (3.22) follows from $-P''(0)/[2P'(0)] = (2n^2 + 1)/3$.

It is worth noting that $\theta_1^{-1} = 2/[1 - \cos(\pi/n)]$ also grows quadratically for large n . A series expansion shows that $d/dn([1 - \cos(\pi/n)](2n^2 + 1)) > 0$ for $n \geq 1$. Thus, the ratio of θ_1^{-1} to $\sum_1^n \theta_k^{-1}$ decreases monotonically from 1 (at $n = 1$) to $6/\pi^2 \approx 0.60793$ (as $n \rightarrow \infty$). This shows that for $n \geq 1$, the cost of computing $\tilde{m}(\lambda_1)$ is more than 60% of the total cost of computing $\tilde{m}(\lambda_k)$ at all the collocation points along the complex line segment Γ —the closest point to the real λ axis dominates the others.

3.3.4. Bound on the extrapolation error. As explained in section 2.4, if $\rho(s)$ is real analytic at $s = \lambda > 0$, then $m(\zeta)$ can be analytically continued from the upper half-plane to a disk of radius a centered at λ . By (3.10), the same is true of $\tilde{m}(\zeta)$. Thus, there exists $\alpha > 1$ such that $\tilde{m}(\zeta)$ is analytic and bounded in the region B bounded by the Bernstein ellipse [45]

$$\zeta = \lambda + i\frac{\ell}{2} \left(1 - \frac{z + z^{-1}}{2} \right), \quad z = \alpha e^{i\theta}, \quad 0 \leq \theta \leq 2\pi,$$

which encloses Γ , has foci at its endpoints, and has semimajor and semiminor axes equal to $\ell(\alpha + \alpha^{-1})/4$ and $\ell(\alpha - \alpha^{-1})/4$, respectively. In particular, it is not hard to show that the choice

$$(3.23) \quad \alpha = \sqrt{1 + \frac{a^2}{\ell^2}} + \frac{a}{\ell} + \sqrt{2} \left(\frac{a}{\ell} \sqrt{1 + \frac{a^2}{\ell^2}} + \frac{a^2}{\ell^2} \right)^{1/2}$$

causes ∂B to cross the real ζ -axis at $\lambda \pm a$ and to remain strictly above the semicircle $\{\zeta : |\zeta - \lambda| = a, \text{Im} \zeta < 0\}$. We regard ℓ as a free parameter of the method and a as a property of the operator L : it is the (effective) radius of convergence of $\rho'(s)$ at $s = \lambda$. Note that a , α , and z here are not related to those in (3.11), (2.12), and (3.21). From (3.23), we see that $\alpha \sim 4a/\ell$ for $\ell \ll a$ and $\alpha \sim 1 + \sqrt{2a/\ell}$ for $\ell \gg a$. The significance of α is that the Chebyshev coefficients of $\tilde{m}(\zeta)|_\Gamma$ decay like α^{-j} ; i.e., we can write

$$(3.24) \quad \tilde{m}(\lambda + i\theta\ell) = \sum_{j=0}^{\infty} a_j T_j(1 - 2\theta), \quad 0 \leq \theta \leq 1, \quad |a_j| \leq 2M\alpha^{-j},$$

where the T_j are Chebyshev polynomials and $M = \max_{\zeta \in \bar{B}} |\tilde{m}(\zeta)|$ (see [45]). Let $p(\theta)$ be the polynomial of degree n that interpolates the value of $\tilde{m}(\lambda + i\theta\ell)$ at $\{\theta_k\}_{k=0}^n$ (including $k = 0$). Then $p(0) = \tilde{m}(\lambda^+)$ and the Chebyshev coefficients of $p(\theta) = \sum_0^n \hat{p}_j T_j(1 - 2\theta)$ satisfy $|\hat{p}_j - a_j| \leq 2M\alpha^{-n}/(\alpha - 1)$ due to aliasing [45]. The polynomial

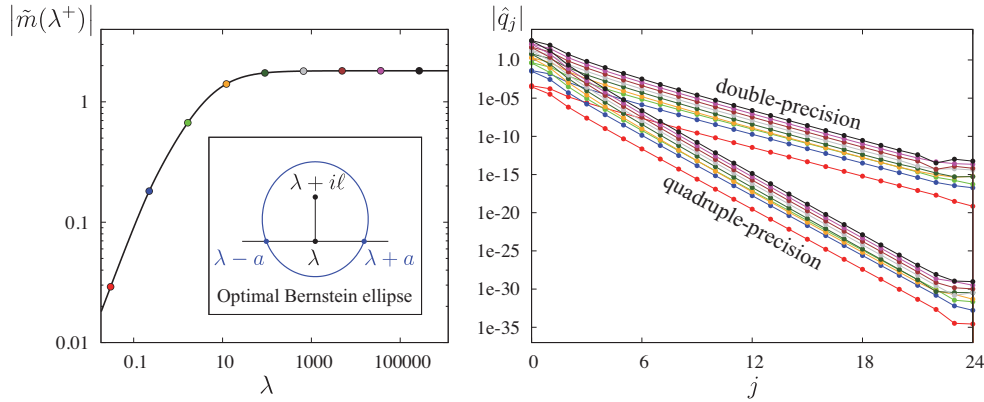


FIG. 4. Plots of $|\tilde{m}(\lambda^+)|$ and the magnitudes of the Chebyshev coefficients \hat{q}_j used to compute $\tilde{m}(\lambda^+)$. The points at the left are color coded to correspond to two curves at the right, one for each type of machine precision used to compute it. The interval lengths ℓ in (3.27) were chosen so the Chebyshev coefficients decay just fast enough to reach machine precision at $j = 24$. The inset at left shows the optimal Bernstein ellipse with $\alpha = 4.9834$.

$q(\theta)$ of degree $n - 1$ that interpolates \tilde{m} at $\{\theta_k\}_{k=1}^n$ differs from $p(\theta)$ by a multiple of the Lagrange polynomial

$$l_0(\theta) = \prod_{j=1}^n \frac{\theta - \theta_j}{0 - \theta_j} = \frac{T_0(1 - 2\theta) + T_n(1 - 2\theta)}{2n} + \sum_{j=1}^{n-1} \frac{T_j(1 - 2\theta)}{n}.$$

The multiple is chosen to reduce the degree of $p(\theta)$ by one, i.e., $q(\theta) = p(\theta) - 2n\hat{p}_n l_0(\theta)$. As a result,

$$(3.25) \quad |q(0) - \tilde{m}(\lambda^+)| = 2n|\hat{p}_n| \leq 2n(|a_n| + |\hat{p}_n - a_n|) \leq \frac{4nM\alpha^{1-n}}{\alpha - 1}.$$

In addition to this bound on the extrapolation error, we find that the Chebyshev coefficients of q agree closely with those of $\tilde{m}(\zeta)|_\Gamma$:

$$(3.26) \quad |\hat{q}_j - a_j| \leq |\hat{q}_j - \hat{p}_j| + |\hat{p}_j - a_j| \leq \frac{6M\alpha^{1-n}}{\alpha - 1} \quad (0 \leq j < n),$$

where we used $|\hat{q}_j - \hat{p}_j| = |2n\hat{p}_n \hat{l}_{0j}|$, which is bounded by $2\hat{p}_n$.

3.3.5. Numerical computation of optimal segment length. Figure 4 confirms that the Chebyshev modes of $q(\theta)$ decay exponentially, as predicted from $|a_j| \leq 2M\alpha^{-j}$ and (3.26). The function $\tilde{m}(\lambda^+)$ computed here corresponds to x^* as defined in section 3.2 rather than the simplified version in (3.19). We find empirically that with $n = 25$ collocation points, the segment lengths

$$(3.27) \quad \ell = \frac{13\lambda}{11 + \lambda^{-11/8}} \text{ (double precision),} \quad \ell = \frac{1.85\lambda}{9 + \lambda^{-11/8}} \text{ (quadruple)}$$

are close to optimal over a wide range of λ values. Indeed, for each λ in the range computed in Figure 4 ($e^{-4} \leq \lambda \leq e^{14}$), the Chebyshev modes \hat{q}_j reach the roundoff threshold just before j reaches $24 = n - 1$. Assuming the observed modes $|\hat{q}_j|$ are a good predictor of the bound $|a_j| \leq 2M\alpha^{-j}$, we can estimate $2M\alpha^{1-n}$ by $|\hat{q}_{n-1}|$. If

the latter is $O(\delta)$, then $|q(0) - \tilde{m}(\lambda^+)|$ is at most $2n/(\alpha - 1)$ times larger, or $O(n\delta)$, due to (3.25). The value of α is substantially larger than 1 since reducing $|a_j|$ by a factor of δ when j varies from 0 to $n - 1$ requires $\alpha \gtrsim \delta^{-1/(n-1)}$. When $n = 25$, this gives $\alpha \gtrsim 4.2$ in double-precision and $\alpha \gtrsim 18$ in quadruple-precision arithmetic. We note from (3.23) that α can be controlled easily by varying ℓ . In practice, we choose ℓ to be just small enough that roundoff errors in the highest-frequency Chebyshev modes of \hat{q}_j begin to be visible, hitting a plateau rather than continuing to decay exponentially.

3.3.6. Roundoff error. The effects of floating-point arithmetic can be taken into account using the extrapolation formula

$$(3.28) \quad q(0) = (-1)^{n-1}q(\theta_n) - 2 \sum_{k=1}^{n-1} (-1)^k q(\theta_k).$$

In exact arithmetic, if $q(\theta)$ interpolates the exact values of $\tilde{m}(\lambda + i\ell\theta_k)$ for $1 \leq k \leq n$, we have shown how to choose ℓ so that the right-hand side of (3.28) will agree with $\tilde{m}(\lambda^+)$ to an error tolerance of $O(n\delta)$. If each term of the right-hand side of (3.28) is perturbed by $O(\delta)$, the accumulated effect remains $O(n\delta)$. In practice, the errors really are this small, as shown in the follow-up paper [48] by comparing double-precision results to “exact” solutions computed in quadruple-precision, and indirectly in [49] by comparing this method of solving (2.4) to a projected dynamics approach using orthogonal polynomials.

3.3.7. Complexity estimate. We conclude this section with a complexity estimate for the algorithm when n (the number of collocation points) and ν (the number of fully implicit Runge–Kutta stages) are chosen optimally.

First, for fixed λ and δ , we wish to find n and ℓ to minimize N_{tot} in (3.22) subject to the constraint $\alpha = \delta^{-1/(n-1)}$. A nearly identical optimization problem (with a simpler answer) is to minimize n^2/ℓ subject to $\alpha = \delta^{-1/n}$, where α is regarded as a function of ℓ (with a fixed) in (3.23). A routine calculation shows that the optimal solution satisfies

$$\alpha(\ell) \ln \alpha(\ell) + 2\ell\alpha'(\ell) = 0, \quad n = \frac{\ln(1/\delta)}{\ln \alpha(\ell)},$$

which yields $\ell/a = 0.9065$, $\alpha = 4.9834$, and $n = 0.6226 \ln(1/\delta)$. This gives $n = 23$ for $\delta = 10^{-16}$ and $n = 46$ for $\delta = 10^{-32}$. Thus, our choice of $n = 25$ above is nearly optimal in double-precision but too small in quadruple-precision, which is what we observe in practice as well (see section 5). Note that the eccentricity of the optimal Bernstein ellipse equals $2/(\alpha + \alpha^{-1}) = 0.3858$, independently of λ and δ (see Figure 4). Its size is determined by a , the effective radius of convergence of $\rho'(s)$ at $s = \lambda$. We can estimate a as a function of λ indirectly, using $\ell/a = 0.9065$ and the left formula in (3.27), since $n = 25$ is close to optimal in double-precision. The target accuracy δ affects the number of grid points via $n = 0.6226 \ln(1/\delta)$, but not the size or shape of the ellipse. With these parameter choices (still approximating $2n^2 + 1$ by $2n^2$), we obtain

$$(3.29) \quad N_{\text{tot}} \leq 3.107 K_\nu(\delta) \ln^3(1/\delta) \frac{\lambda}{\ell},$$

where $\lambda/\ell \approx 11/13 + (1/13)\lambda^{-11/8}$ in our case, and $K_\nu(\delta) = \max\left(\frac{3}{2}, \frac{4}{9}\delta^{-1/(2\nu-1)}\right)$.

The additional $(77/6) \max(\lambda^{-2}, \lambda^{1/2}) K_\nu(\delta)$ steps required to evolve from 0 to x^* are normally a small fraction of N_{tot} . These initial steps do not require complex arithmetic, which further reduces their cost. Nevertheless, λ^{-2} will eventually dominate $\lambda^{-11/8}$ when $\lambda \rightarrow 0$. Over the range of λ and tolerances considered here ($e^{-4} \leq \lambda \leq e^{14}$, $\delta \leq 10^{-15}$), $3.107 \ln(1/\delta)^3$ is at least nine times larger than $(77/6)\lambda^{1/2}$. But for still larger values of λ , the latter could become significant. Fortunately, this term is an artifact of choosing $x^* = 1$ in (3.19) and is not present when x^* is defined as the first negative extremum of $y_1(x; \lambda)$. The naive algorithm (in which x^* is set to zero and only one fundamental matrix is computed) can even be used since there is no growth phase when $\lambda \geq V_{\text{max}} = 0.18704$.

The difficulty in analyzing the case in which λ is large and x^* is allowed to drop below 1 is that the estimate (3.16) of the convergence rate is based on the large x asymptotics of $\Psi(x)$, yielding errors of the form $\exp(-\gamma[x^{5/2} - (x^*)^{5/2}])$, where γ is a multiple of $\text{Im} \sqrt{\lambda_k}$. This formula (incorrectly) suggests that if x^* is close to zero, little progress will be made until x reaches 1. By contrast, (3.18) does not rely on asymptotics and shows that the number of steps needed to evolve from x^* to x scales linearly, like $|\lambda_k|^{1/2}[x - x^*]$, when x^* is close to zero and $x \leq 1$. We expect that a more refined WKB analysis using (C.13) instead of (C.15) would show that for large λ , $|\tilde{y}_0/\tilde{y}_1 + \tilde{m}|$ decays initially like $\exp(-\tilde{\gamma}[x - x^*])$, where $\tilde{\gamma}$ is also a multiple of $\text{Im} \sqrt{\lambda_k}$, before exhibiting the $\exp(-\gamma[x^{5/2} - (x^*)^{5/2}])$ behavior. As a basic check, if we approximate $\sqrt{F(x)/\Psi(x)}$ in (C.13) by $\sqrt{F(0)/\Psi(0)}$ and assume that $h_\pm(x)$ in (C.16) remain small down to $x = 0$ when $\lambda \gg 1$, then $\epsilon(x)$ in (3.17) becomes $\exp(-\tilde{\gamma}[x - x^*])$ with $\tilde{\gamma} = \sqrt{6}\pi^{1/4} \text{Im}\{\sqrt{\lambda_k}\} + O(|\lambda_k|^{-1/2})$. Since the number of steps and the logarithm of the inverse error grow similarly as functions of x , with prefactors proportional to $|\lambda_k|^{1/2}$ and $\text{Im}\{\sqrt{\lambda_k}\}$, respectively, the technique of linking N_{steps} to $\ln(1/\delta)$ by comparing (3.20) to (3.21) should work the same. We did not carry out the details as this was not an issue over the range of λ considered here.

Next, to leading order in ν , the computational cost of the steps in (3.29) is $C_{\text{tot}} = C(\delta) \times \frac{1}{3}(2\nu)^3 \times 8 \times N_{\text{tot}}$, where $C(\delta)$ is the cost of a floating-point operation with roundoff threshold δ , $\frac{1}{3}(2\nu)^3$ is the number of multiplications required to solve the $2\nu \times 2\nu$ linear system associated with a ν -stage fully implicit RK step for a linear ODE with two components (see [48]), and 8 accounts for the four multiplications and four additions required to perform one complex multiplication and one complex addition. Optimizing C_{tot} boils down to minimizing $\nu^3 \delta^{-1/(2\nu-1)}$ with δ fixed. Minimizing $(2\nu - 1)^3 \delta^{-1/(2\nu-1)}$ instead gives $2\nu - 1 = \frac{1}{3} \ln(1/\delta)$, for which $K_\nu(\delta) = \frac{4}{9} \exp(3) \approx 8.927$. Using $2\nu \leq \frac{10}{9}(2\nu - 1)$ for $\nu \geq 5$, we obtain

$$(3.30) \quad C_{\text{tot}} \leq 3.76 C(\delta) \ln^6(1/\delta) \frac{\lambda}{\ell}.$$

Since $C(\delta)$ grows like $\ln^2(1/\delta)$ or $\ln(1/\delta) \ln[\ln(1/\delta)]$, depending on the arbitrary precision implementation, we conclude that the cost of computing $\rho'(\lambda)$ by this method with accuracy δ grows more slowly than $\delta^{-\gamma}$ for any $\gamma > 0$ as $\delta \rightarrow 0$. This translates into significant performance gains over other methods for computing spectral density functions with high accuracy, as shown in section 5 and [48].

We find that the estimate $n = 0.6226 \ln(1/\delta)$ for the optimal number of collocation points agrees closely with comparisons of actual running times in practice—the same n is close to optimal. However, the bounds (3.20) and (3.29) overpredict the number of steps taken using adaptive stepsize control [48] by a factor of roughly $3(2\nu - 1)$, and the optimal choice of $(2\nu - 1)$ is a few times larger than $\frac{1}{3} \ln(1/\delta)$. This factor

of $(2\nu - 1)$ can be obtained in the analysis of Appendix A if we set $rM = (2\nu - 1)$ instead of $5/3$ in (A.2). (M is a type of Lipschitz constant for the ODE, r is the radius of a disk centered at x in the complex plane, and h is the stepsize.) However, for technical reasons explained in [48], our current analysis requires $hM \leq 2/3$, which breaks down if r is increased in this way. We believe it should be possible to remove this barrier using A-stability of the scheme rather than a Neumann series to bound the condition number of the IRK equations, but we do not know how to do this. If it is indeed possible to increase rM to $(2\nu - 1)$ without losing control of the constants in the formulas, then $K_\nu(\delta)$ would be of the form $\frac{\text{const}}{2\nu-1}\delta^{-1/(2\nu-1)}$, and the result (3.30) could be improved to contain $\ln^5(1/\delta)$ instead of $\ln^6(1/\delta)$.

We remark that these optimization problems are only intended to serve as a guideline for choosing n and ν and deriving a rough complexity estimate. The correctness of the algorithm does not depend on choosing n or ν optimally and does not rely on asymptotics.

4. Numerical examples. We now consider two examples illustrating the use of (2.30), (2.31) to solve $u_t = -Lu$ with initial conditions $u(x, 0) = f_j(x)$:

$$(4.1) \quad \text{Example 1: } f_1(x) = x. \quad \text{Example 2: } f_2(x) = x^2.$$

Example 1 is harder to compute since Lf_1 has a singularity at $x = 0$ that leads to an infinite initial speed u_t there. This causes $\hat{f}_1(\lambda)$ to decay slowly (like λ^{-2}), just as the Fourier transform of a function with a slope-discontinuity decays slowly. Nevertheless, this singular example is relevant to the problem of resistive damping in a plasma; see [29, section 3].

We computed $\hat{f}_1(\lambda)$ and $\hat{f}_2(\lambda)$ at the grid points λ_j in (3.7) using the method explained above, in which a third component is added to \vec{r} in (2.12) to represent $F(x) = \int_0^x se^{-s^2/2}f(s)y_1(s; \lambda) ds$, and the solution is evolved until $F(x)$ reaches its limiting value. The results are shown in Figure 5, where we have adopted the notation

$$(4.2) \quad \begin{aligned} \tilde{f}(\sigma, t) &:= \hat{f}(e^\sigma) \exp(-e^\sigma t) Y(e^\sigma) \rho'(e^\sigma) e^\sigma, & \tilde{u}(x, t) &= \int_{-\infty}^{\infty} \frac{y_1(x; e^\sigma)}{xY(e^\sigma)} \tilde{f}(\sigma, t) d\sigma. \\ \tilde{u}(x, t) &:= [u(x, t) - 4\hat{f}(0)/\sqrt{\pi}] e^{-x^2/2}, \end{aligned}$$

The tilde here is not related to the one used in (3.9)–(3.12) to denote solutions of (2.12) starting at $x = x^*$. Note that $\tilde{f}(\sigma, t) = \hat{u}(\lambda, t) Y(\lambda) \rho'(\lambda) \lambda$ with $\hat{u}(\lambda, t) = \hat{f}(\lambda) e^{-\lambda t}$ and $\lambda = e^\sigma$, and the extra factor of λ accounts for $d\lambda = e^\sigma d\sigma$ in the change of variables. Also, if $\tilde{f}(\sigma, t^*)$ can be represented efficiently as a function of σ for some fixed $t^* \geq 0$, then for $t \geq t^*$ we have

$$(4.3) \quad \tilde{f}(\sigma, t) = \tilde{f}(\sigma, t^*) e^{-e^\sigma(t-t^*)},$$

which is easy to evaluate. We will represent $\tilde{f}(\sigma, t^*)$ using a Fourier series on $-4 = \sigma_L \leq \sigma \leq \sigma_R = 14$, where $t^* = 10^{-4}$ in Example 1 and $t^* = 0$ in Example 2.

Figure 6 shows that the grid spacing in (3.7) is sufficient to represent $\tilde{f}_1(\sigma, t)$ over $-4 \leq \sigma \leq 14$ for $t \geq 10^{-4}$ and $\tilde{f}_2(\sigma, t)$ for $t \geq 0$ to quadruple-precision accuracy using the FFT. The red \times markers were obtained by computing $\hat{f}_2(\lambda)$ and $\rho'(\lambda)$ directly. The black markers were obtained from the red by truncating the data at $\sigma = 6.921875$ (the vertical red line), applying the FFT (to all 768 points), truncating the Fourier series at $k = 317$ (the vertical green line), and transforming back. The blue markers show a filtered version of the raw data labeled $t = 10^{-4}$ in Figure 5. In this case, the




FIG. 5. Plots of $|\hat{f}(\lambda)|$ and $|\tilde{f}(\sigma, t)|$ at various times for the two examples in (4.1). The effective support (exceeding 10^{-30}) of $\tilde{f}_2(\sigma, t)$ lies between $-4 < \sigma < 7$ for $t \geq 0$, while that of $\tilde{f}_1(\sigma, t)$ extends beyond $\sigma = 14$ until $t = 10^{-4}$ due to the slow λ^{-2} decay rate of $\hat{f}_1(\lambda)$. Both $\tilde{f}_1(\sigma, t)$ and $\tilde{f}_2(\sigma, t)$ turn out to be negative everywhere they can be distinguished from roundoff error.




FIG. 6. Sampling $\tilde{f}_1(\sigma, 0.0001)$ and $\tilde{f}_2(\sigma, 0)$ at the 768 grid points in (3.7) is sufficient to reduce the highest frequency Fourier modes to roundoff error in quadruple-precision arithmetic. (Right) The raw data (red) is evaluated on the original grid, while the filtered data (black and blue) is evaluated on a finer mesh with 1536 grid points.

Fourier series was truncated at $k = 317$ with no initial filter in σ . Note that roundoff error causes $|\tilde{f}(\sigma, t)|$ to grow to around 10^{-30} near $\sigma = -4$ in both examples. This is not a problem since $Y(\lambda)$ was included in the raw data before the FFT was computed. For larger values of t in Example 1 and all $t \geq 0$ in Example 2, the domain can be reduced to $-4 \leq \sigma \leq 8$ so that only 512 raw data points are needed. However, we will continue to work with the grid (3.7) for illustration.

To perform the integral in (4.2), we use the trapezoidal rule over the interval $-4 \leq \sigma \leq 14$. We increase the number of collocation points as x increases in order to resolve the increasingly oscillatory integrals involved. Much of this work can be done once and for all, independent of the initial condition $f(x)$. To this end, we precompute

$$(4.4) \quad g(x; \lambda) = \frac{y_1(x; \lambda)}{xY(\lambda)}, \quad \lambda = e^\sigma,$$

at selected x locations on a nested hierarchy of grids

$$(4.5) \quad \sigma_j^{(p)} = -4 + \frac{3j}{128 \times 2^p}, \quad 0 \leq j < 768 \times 2^p, \quad 0 \leq p \leq 10.$$

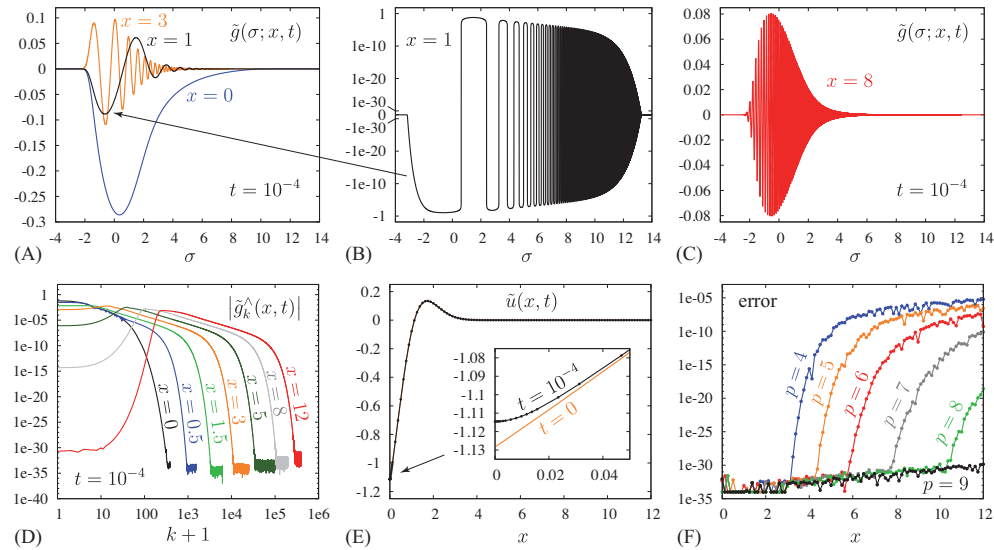


FIG. 7. Plots of $\tilde{g}(\sigma; x, t)$, its Fourier modes $\tilde{g}_k^\wedge(x, t)$, the solution $\tilde{u}(x, t)$, and the error $|\tilde{u}^{(p)}(x, t) - \tilde{u}^{(10)}(x, t)|$ for Example 1 at $t = 10^{-4}$. As x increases, $\tilde{g}(\sigma; x, t)$ becomes more oscillatory, and more grid points are needed to avoid aliasing errors in the trapezoidal rule integration scheme.

For $p \geq 1$, $g(x, \exp(\sigma_j^{(p)}))$ has to be computed only for odd indices j since it is already known for even indices from the previous level. Note that $\sigma_j^{(0)}$ coincides with σ_j in (3.7). This makes it easy to interpolate the values of $\tilde{f}(\sigma, t)$ rather than computing $m(\lambda^+)$ and $\hat{f}(\lambda)$ at the new grid points directly. To increase the size of \tilde{f} by a factor of 2^p , we simply zero-pad its FFT and compute the inverse transform. For example, the black and blue markers in the right panel of Figure 6 were computed in this way on the $p = 1$ grid with 1536 nodes.

The results of this calculation for Example 1 are given in Figure 7. Panels (A)–(C) show the integrand $\tilde{g}(\sigma; x, t)$ in the reconstruction formula

$$(4.6) \quad \tilde{u}(x, t) = \int_{-\infty}^{\infty} \tilde{g}(\sigma; x, t) d\sigma, \quad \tilde{g}(\sigma; x, t) = g(x, e^\sigma) \tilde{f}(\sigma, t),$$

as a function of σ for $x = 0$, $x = 1$, $x = 3$, and $x = 8$ with $t = 0.0001$ fixed. Note that as x increases, \tilde{g} becomes more oscillatory as a function of σ . Panel (B) shows the $x = 1$ solution stretched vertically to a signed logarithmic scale. This was done by plotting $\arcsinh(\frac{1}{2}10^{32}\tilde{g})$ on the y -axis and placing tick marks where $\tilde{g} = \pm 10^{-10k}$. Note that \tilde{g} becomes highly oscillatory as it decays. On a stretched scale, solutions with other values of x have a similar envelope to the $x = 1$ solution shown here, and range from having no oscillations ($x = 0$) to very rapid oscillations ($x = 12$). Panel (D) gives the magnitudes of the Fourier modes of $\tilde{g}(\sigma; x, t)$ with x and t held fixed. The mode amplitudes of the FFT are normalized by $1/N$, where $N = 768 \times 2^p$ is the number of grid points. With this scaling, the N -point trapezoidal rule gives $\tilde{u}(x, t) = 18\tilde{g}_0^\wedge(x, t)$, where $18 = \sigma_{\max} - \sigma_{\min}$. The curves labeled $x = 0$, $x = 0.5$, etc., were computed with $p = 0, 2, 4, 6, 8, 9$, and 10 , respectively. These levels were chosen so that \tilde{g}_k^\wedge decays to roundoff error before k reaches the Nyquist frequency $k = N/2$, which is the largest mode shown for each curve. Panel (E) shows the solution $\tilde{u}(x, t)$

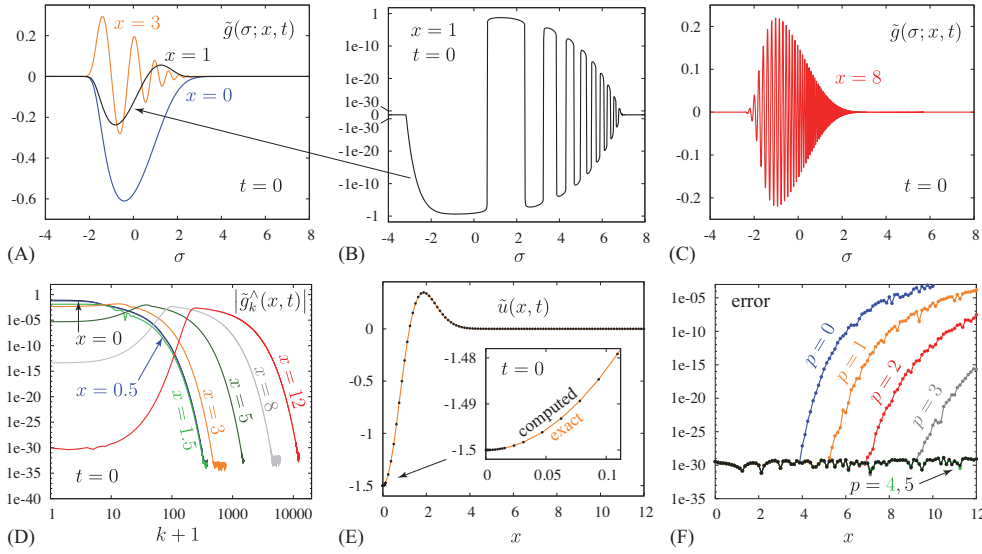


FIG. 8. Analogous results for Example 2. Note that the axes are scaled differently in Figure 7.

at $t = 10^{-4}$, obtained by integrating $\tilde{g}(\sigma; x, t)$. For comparison, we also plot

$$(4.7) \quad \tilde{u}(x, 0) = \left(x - \frac{2}{\sqrt{\pi}}\right) e^{-x^2/2},$$

which agrees closely with $\tilde{u}(x, 10^{-4})$ except near $x = 0$, where u_t is initially infinite. Panel (F) gives the error in the reconstructed solution using $p = 10$ as the exact solution. Higher values of x require finer grids to resolve the oscillations in $\tilde{g}(\sigma; x, t)$. For this example, the $p = 9$ and 10 solutions are identical to 30 digits of accuracy.

In Figure 8, we present analogous results for Example 2. Since L maps $f(x) = x^2$ to a smooth function that is well behaved at the origin, the reconstruction can be done at $t = 0$ to recover $u(x, 0) = f(x)$. By contrast, we needed $t \geq 10^{-4}$ to overcome the slow decay of $\hat{f}(\lambda)$ in Example 1. The main difference between Figures 7 and 8 is that the effective support of \tilde{g} is smaller in the latter case, leading to less oscillatory integrals. This is evident on comparing panel (B) in both figures, and also on observing that fewer Fourier modes are needed in panel (D) to reach machine-precision in Example 2. More specifically, the curves labeled $x = 0, x = 0.5$, etc., in (D) were computed with $p = 0, 0, 0, 1, 2, 4$, and 5 , respectively. In panel (E), the orange curve gives the exact initial condition

$$(4.8) \quad \tilde{u}(x, 0) = \left(x - \frac{3}{2}\right) e^{-x^2/2},$$

while the black markers are computed using the trapezoidal rule on the numerically computed $\tilde{g}(\sigma; x, t)$. Panel (F) gives the true errors relative to the exact initial condition (4.8) rather than using the solution on the finest grid as the benchmark. In Example 2, the solution reaches roundoff error already at $p = 4$, which corresponds to a grid 32 times coarser than the $p = 9$ mesh needed in Example 1.

The accelerating frequency of oscillation that occurs in $\tilde{g}(\sigma; x, t)$ as σ increases is partly due to our choice of $\lambda = e^\sigma$ in the change of variables (4.2). When x is small,

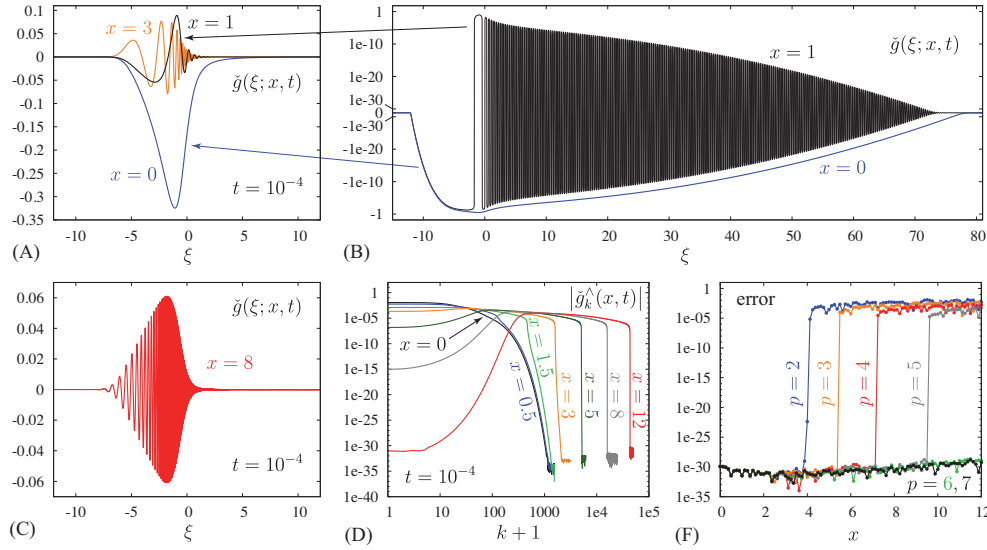


FIG. 9. The alternative change of variables (4.9) leads to a more efficient representation of the oscillatory integrand $\tilde{g}(\xi; x, t)$ for $x \geq 5/8$. Panel (E) is omitted as it looks identical to Figure 7(E).

this choice is very good for representing $\tilde{g}(\sigma; x, t)$ with a small number of Fourier modes. However, we can do better for larger x . From the asymptotic analysis in Appendix C, we expect $x^{-1}y_1(x; \lambda)$ in (4.4) to oscillate like $x^{-1/4} \cos(\sqrt{8}x^{5/2}\sqrt{\lambda})$ at leading order. Thus, to achieve a nearly constant number of grid points per cycle with respect to λ holding x fixed, we should change variables so that $\lambda \sim \xi^2$ for large ξ . We also want $\lambda \rightarrow 0$ as $\xi \rightarrow -\infty$. We tried functions of the form

$$(4.9) \quad \lambda = A(\sqrt{1 + \xi^2} + \xi)^2 = A(\sqrt{1 + \xi^2} - \xi)^{-2}$$

and found that $A = 25$ works nicely. The first formula is used for positive ξ and the second for negative ξ .

Figure 9 shows the results for Example 1 at $t = 0.0001$ with the alternative integration variable. Plots (A)–(C) show the integrand $\tilde{g}(\xi; x, t)$ in the reconstruction formula

$$(4.10) \quad \tilde{u}(x, t) = \int \tilde{g}(\xi; x, t) d\xi, \quad \tilde{g}(\xi; x, t) = g(x; \lambda)\check{f}(\xi, t),$$

where λ is related to ξ via (4.9), $\check{f}(\xi, t) = 2\tilde{f}(\sigma, t)/\sqrt{1 + \xi^2}$, and $\sigma = \ln \lambda$. Note that the oscillation frequency is nearly uniform over $0 \leq \xi \leq 81$ in (B), unlike the result in Figure 7. To compute the integrals, we precompute $g(x; \lambda)$ on a nested grid similar to (4.5) but over $-15 \leq \xi \leq 81$, namely,

$$(4.11) \quad \xi_j^{(p)} = -15 + \frac{j}{8 \times 2^p}, \quad 0 \leq j < 768 \times 2^p, \quad 2 \leq p \leq 7.$$

The lowest level is labeled 2 so that the number of grid points at a given level is the same in Figures 7–9. The drawback of using ξ is that more grid points are needed at $x = 0$ to represent $\tilde{g}(\xi; x, t)$ than $\tilde{g}(\sigma; x, t)$. The benefit is that fewer grid points are needed for larger x . The following table gives the index N at which the Fourier modes

of $\tilde{g}(\sigma; x, t)$ and $\tilde{g}(\xi; x, t)$ reach roundoff error in quadruple-precision with $t = 0.0001$:

x	0	5/8	1	3	8	12
N_σ	320	1200	2000	11420	109000	290000
N_ξ	1160	1220	1300	2120	16000	44000

We use these numbers as a guideline for the optimal number of collocation points to use in the trapezoidal rule. As x increases, the cost of reconstructing the solution via (4.6) or (4.10) increases due to more collocation points being needed to resolve the oscillations in \tilde{g} or \tilde{g} and more timesteps being needed to evolve to x . Note that the σ variable is better for small x because \tilde{g} and \tilde{g} grow at similar rates near $\sigma = -2$ and $\xi = -5$, respectively, but the domain for ξ is several times larger than for σ . However, for larger x , the oscillations dominate the smoothness properties of \tilde{g} and \tilde{g} and are spread out more uniformly in the ξ variable. Thus, fewer collocation points are wasted in less oscillatory regions. In our code, we use the Fourier representation of $\tilde{f}(\sigma, t)$ in Figure 6 to evaluate $\tilde{f}(\xi, t)$ on the $p = 2$ grid since this requires only 768 evaluations of $m(\lambda^+)$ and $\hat{f}(\lambda)$ rather than 3072. Once $\tilde{f}(\xi, t)$ is known at $p = 2$, we use the FFT to interpolate to higher levels.

5. Summary of the algorithm and comparison with other methods. We give here a brief summary of the method presented in detail in the sections above, which may help readers interested in implementing it. The key steps of the algorithm may be summarized as follows:

1. Pick a “coarse” mesh in λ -space, e.g., $\lambda_j = e^{\sigma_j}$ or $25[(1 + \xi_j^2)^{1/2} - \xi_j]^{-2}$, where σ_j (or ξ_j) is uniformly spaced from σ_L to σ_R (or ξ_L to ξ_R). The number of grid points and choice of σ_L and σ_R may be adjusted later, in steps 5 and 6.
2. For each λ on the mesh, evolve $y_1(x; \lambda)$ forward to x^* , the location of its first negative extremum, $y_{\max}(\lambda)$. If the ODE is singular at the origin, i.e., (2.12) is used rather than (2.16), use the series solution (2.14) to initialize the ODE to the right of the singularity, e.g., at $x = 10^{-6}$. Use enough terms of the series to achieve roundoff-level accuracy. Also define $Y(\lambda)$ as in (3.8) and evolve $y_0(x; \lambda)$ to x^* to obtain $\Phi(x^*; \lambda)$ in (3.9). We use an arbitrary (e.g., 50th) order, fully implicit Runge–Kutta collocation (IRK) method to advance the solutions in x .
3. Let $\lambda_k = \lambda + i\ell\theta_k$, with λ a grid point from steps 1 and 2 and $\theta_1, \dots, \theta_n$ the Chebyshev points in (3.2). Evolve $\tilde{\Phi}(x; \lambda_k)$ forward from the identity at $x = x^*$ until its columns become linearly dependent to machine precision (determined by monitoring the Wronskian). Record $\tilde{m}(\lambda_k) = -\tilde{y}_0(x; \lambda_k)/\tilde{y}_1(x; \lambda_k)$. Extrapolate to obtain $\tilde{m}(\lambda^+)$ as well as $m(\lambda^+)$ and $\rho'(\lambda)$. Adjust ℓ , if necessary, to obtain appropriate Chebyshev mode decay rates (Figure 4).
4. Compute the transform of the initial condition at the grid points via $\hat{f}(\lambda) = \lim_{x \rightarrow \infty} F(x)$, $F(x) = \int_0^x [w^{1/2}(s)f(s)]y_1(s; \lambda) ds$. F is evolved via an ODE, along with y_1 and z_1 in (2.11), until F stops changing, which happens rapidly due to $w^{1/2}(x) = xe^{-x^2/2}$. If the ODE is singular at the origin, initialize F at $x = 10^{-6}$ using the series solution of the ODE.
5. Evaluate $\tilde{f}(\sigma, t^*) = \hat{f}(e^\sigma) \exp(-e^\sigma t^*) Y(e^\sigma) \rho'(e^\sigma) e^\sigma$ on the grid at $t^* = 0$. If necessary, go back to step 1, and adjust the mesh endpoints so that \tilde{f} decays rapidly to zero as $\sigma \rightarrow \sigma_L^+$ and $\sigma \rightarrow \sigma_R^-$. It may be necessary to increase t^* to achieve sufficient decay at the right endpoint (Figure 5).

6. Compute the FFT of $\tilde{f}(\sigma, t^*)$ for the purpose of interpolation. Adjust the number of mesh points (in step 1) as necessary so the Fourier modes of $\tilde{f}(\sigma, t^*)$ decay to roundoff accuracy. Optionally, filter the modes (Figure 6).
7. Compute the basis functions $g(x; \lambda) = y_1(x; \lambda)/[xY(\lambda)]$ by solving the ODE for y_1 on a nested hierarchy of grids, $\lambda = \exp\{\sigma_j^{(p)}\}$, as in (4.5). Record $g(x; \lambda)$ at the points x where $u(x, t)$ is to be evaluated. Start with level $p = 0$ and add levels as needed in step 8.
8. Compute $\tilde{g}(\sigma; x, t) = g(x, e^\sigma)\tilde{f}(\sigma, t)$ on successive levels of the grid hierarchy and evaluate $u(x, t) = \int_{\sigma_L}^{\sigma_R} \tilde{g}(\sigma; x, t) d\sigma$ via the trapezoidal rule. Since the trapezoidal rule is the constant mode of the FFT, stop at level p of the grid hierarchy when the FFT of $\tilde{g}(\sigma; x, t)$ decays to roundoff accuracy. Evaluation of $\tilde{f}(\sigma, t)$ is done via (4.3), with $\tilde{f}(\sigma, t^*)$ evaluated via the FFT from step 6.

Steps 5–8 should be modified via (4.10) and (4.11) if ξ is used instead of σ in step 1. We note that the procedure is very general: L can be replaced by any singular Sturm–Liouville operator that is regular (or of limit circle type, with Taylor or Frobenius series solutions) at the origin and of limit point type at ∞ , provided the spectral density function $\rho'(\lambda)$ is smooth enough. Smoothness of ρ affects the smoothness of $m(\lambda)$ in (2.22) and (2.27), which determines how effectively $m(\lambda)$ can be extrapolated to the real axis from the upper half-plane via Chebyshev polynomials. Steps 5–8 also rely on ρ being sufficiently smooth to ensure that the Fourier modes of \tilde{f} and \tilde{g} decay rapidly. We note that each step of the algorithm provides a posteriori error estimates based on the decay of Chebyshev and Fourier modes. This allows rapid selection of mesh parameters through a few iterations of steps 1, 5, and 6.

Since the new algorithm is based on an exact mathematical formula expressing the solution $u(x, t)$ at a later time in terms of the initial condition $f(x)$, it will be as accurate as the quantities $\hat{f}(\lambda)$, $y_1(x; \lambda)/[xY(\lambda)]$, and $[Y(\lambda)\rho'(\lambda)]$ that appear in (2.30), up to quadrature error in the trapezoidal rule, which is controlled by checking that the Fourier modes decay to machine precision. Computing $\hat{f}(\lambda)$, $y_1(x; \lambda)$, and $Y(\lambda)$ involves solving linear ODEs, which are easily solved to machine precision using high order Runge–Kutta methods. Achieving high relative accuracy in $y_1(x; \lambda)$ requires that it not be a multiple of the recessive solution, but this is guaranteed for $\lambda \neq 0$ since the point spectrum of L is $\{0\}$. Computing $m(\lambda) = -\lim_{x \rightarrow \infty} \frac{y_0(x; \lambda)}{y_1(x; \lambda)}$ can also be done to machine precision using high order Runge–Kutta methods. The extrapolation procedure to compute $m(\tau^+) = \lim_{\varepsilon \rightarrow 0^+} m(\tau + i\varepsilon)$ is numerically stable in finite precision arithmetic, as shown in section 3.3.6. However, high relative accuracy of the complex number $m(\tau^+)$ does not imply high relative accuracy of the imaginary part, and digits can be lost when evaluating $\rho'(\tau) = \frac{1}{\pi} \text{Im}\{m(\tau^+)\}$ when $\tau \ll 1$. We showed how to avoid this loss of accuracy in the imaginary part by factoring the fundamental matrix as in (3.9) and using the Wronskian identity to correct the smaller singular value of one of the factors.

To confirm that the new algorithm is spectrally accurate, we have checked in joint work with Landreman [49] that the solution $u(x, t)$ computed as above agrees to roundoff accuracy (14 or 29 digits) with the projected dynamics of (1.1) in spaces of orthogonal polynomials. Such agreement provides strong evidence that monitoring Chebyshev and Fourier mode decay rates provides accurate a posteriori error estimates. The current algorithm may be viewed as a method of approximating exact integral formulas for the solution, while that in [49] may be regarded as a (nearly) exact evolution of a finite-dimensional approximation of the PDE.

To gauge the performance of the new algorithm, in follow-up work [48], we compare our method of computing the spectral density function (steps 2 and 3) to the popular software package SLEDGE [36, 19, 18] and to the algorithm of Fulton, Pearson, and Pruess (FPP) [15, 16]. The algorithm in SLEDGE is based on the Levitan–Levinson formula, $\rho(\lambda) = \lim_{b \rightarrow \infty} \rho_b(\lambda)$, where ρ_b is the spectral function associated with the regular problem on the interval $0 < x < b$. SLEDGE could not handle the weight function $w(x) = x^2 e^{-x^2}$ due to underflow for $x > 26$ in double-precision. To convert to a constant weight function, we made the change of variables $u = w^{-1/2}y$ to obtain (2.16). In this form, the problem is regular at the origin, as discussed in section 2.2, and we were able to compute the spectral function $\rho(\lambda)$ using SLEDGE to around four digits of accuracy. Further refinement of the mesh led to failure of the algorithm, apparently due to overflow when scaling the eigenfunctions with small eigenvalues to satisfy $\Psi y'_1 = 1$ at $x = 0$. Computing $\rho'(\lambda)$ from $\rho(\lambda)$ would lead to additional loss of accuracy. Thus, SLEDGE was not found to be suitable for our purposes.

The FPP algorithm [15, 16] proceeds by defining a sequence of auxiliary functions P_n, Z_n, R_n such that

$$(5.1) \quad f_n(s; \lambda) = \frac{1/\pi}{P_n(s; \lambda)U_1(s; \lambda)^2 + Z_n(s; \lambda)U_1(s; \lambda)U'_1(s; \lambda) + R_n(s; \lambda)U'_1(s; \lambda)^2}$$

converges to $\rho'(\lambda)$ as $s \rightarrow \infty$, with improved convergence rates as n increases. Here $U_1(s; \lambda)$ is related to $u_1(x; \lambda)$ or $y_1(x; \lambda)$ in previous sections by the Liouville transformation [7, 48] to Schrödinger form, $-U'' + Q(s)U = \lambda U$. For $n = 1$ and $n = 2$, the FPP procedure gives

$$(5.2) \quad P_1 = \sqrt{\lambda}, \quad Z_1 = 0, \quad R_1 = \lambda^{-1/2},$$

$$(5.3) \quad P_2 = \sqrt{\lambda - Q}, \quad Z_2 = -\frac{1}{2}Q' / (\lambda - Q)^{3/2}, \quad R_2 = 1/\sqrt{\lambda - Q}.$$

The formulas for $n = 3$ are given in [15] and [48] along with explicit error estimates. These estimates imply that, in the present case, the error in approximating $\rho'(\lambda)$ by $f_n(s; \lambda)$ decays like $s^{-(10n-8)/5}$, where the Liouville transformation relates x to s via the ODE $dx/ds = \sqrt{\Psi(x)}$; see Lemma B.4, Appendix C, and [48].

The following table gives running times (in seconds) for computing the spectral density function with our method and the FPP method at the 768 grid points $\lambda_j = e^{\sigma_j}$ in (3.7) using a 3.33 GHz Intel Xeon X5680 system with 12 cores:

	FPP4d3	FPP15d3	WC50d25		FPP31q3	WC76q25	WC76q48
time	10637	9.75	16.2		21284	2021	1279

Here the letters d and q stand for double- and quadruple-precision computations with tolerances set to 10^{-15} and 10^{-30} , respectively. FPP4d3 employs the fourth-order timestepper described in [15], which uses Richardson extrapolation to improve the accuracy of a second-order frozen coefficient method, with $n = 3$ in (5.1). FPP15d3 and FPP31q3 also use $n = 3$ in (5.1), but with a 15th- or 31st-order timestepper. The letters WC refer to the new algorithm presented in this article, and WC76q25 means that we use a 38-stage, 76th-order fully implicit RK method with 25 Chebyshev extrapolation points. Comparing WC76q25 to WC76q48 confirms the prediction in section 3.3.7 that $n = 46$ extrapolation points should be close to optimal in quadruple-precision. (Our FFT implementation required adjusting to $n = 48$.) The

transformation to Liouville normal form for the FPP method involves solving the nonlinear system $dx/ds = \sqrt{\Psi(x(s))}$, $dU/ds = W$, $dW/ds = [Q(s) - \lambda]U$, where $x(s)$ is needed to compute $Q(s)$. Since fully implicit methods are difficult to implement for nonlinear equations, we used a spectral deferred correction scheme [12] in the FPP15d3 and FPP31q3 cases. The scheme orders in the table were chosen optimally by trial and error in all but the FPP4d3 cases.

The running times above show that the choice of timestepper is critical for efficiently computing spectral density functions with high accuracy. In double-precision, our method is 650 times faster than the FPP algorithm described in [15]. However, we were able to improve their algorithm to be 1.66 times faster than ours by implementing a better timestepper. In quadruple-precision, our method is 16.6 times faster than their method using the best timestepper available for each method. To explain this, we recall [15, 48] that the error $|f_n(s; \lambda) - \rho'(\lambda)|$ in their method decays like $O(s^{-22/5})$ when $n = 3$. Thus, to reduce the error to $O(\delta)$, the solution must be evolved to $s_{\max} = O(1/\delta)^{5/22}$. Empirically, $s_{\max} \approx 0.5(1/\delta)^{5/22} \lambda^{-7/10}$ works well in double- and quadruple-precision arithmetic over $e^{-4} \leq \lambda \leq e^{14}$, though establishing precise dependence on λ is difficult. Indeed, for smaller values of λ this will not be adequate to traverse the growth region of Figure 2. (The potential $Q(s)$ must drop below λ for the convergence theory of [15] to be valid.) From (3.20) and Lemma B.4, this translates into $N_{\text{steps}} \approx 6K_\nu(\delta)\lambda^{-1/5}(1/\delta)^{5/22}$. By contrast, our method requires $N_{\text{tot}} \approx 0.24K_\nu(\delta)(11 + \lambda^{-11/8})\ln^3(1/\delta)$ steps over this range of λ , where again the λ -dependence is empirical and not precisely known. Since $(1/\delta)^{5/22}$ exceeds $\ln^3(1/\delta)$ once $\delta < 10^{-23}$, this is roughly the transition point where our method should become more efficient. In practice, as shown in the table above, the methods are already comparable in double-precision ($\delta = 10^{-15}$) since converting to Liouville normal form in the FPP method introduces the nonlinear equation $dx/ds = \sqrt{\Psi(x)}$. A more detailed comparison of the convergence rates of the two methods will be given in [48].

An advantage of the FPP method is that λ is real, so the solution remains real, and complex linear algebra is avoided. This comes at the cost of having to evaluate higher derivatives of $Q(s)$ that appear in the formulas for P_n , Z_n , and R_n in (5.2) for $n \geq 3$; see [15, 48] for details. This is a significant issue when $Q(s)$ is complicated and could lead to loss of accuracy if the derivatives are computed numerically. In our case, $Q(s)$ is the term in parentheses in (C.6) below, with x replaced by $x(s)$. In [48], to make this practical, we had to resort to asymptotics, making use of Lemmas B.1 and B.4 from Appendix B below to evaluate $Q'''(s)$ in Z_3 . By contrast, our method involves only the original ODE. No derivatives of the potential need to be computed, and the equation need not even be converted to Liouville normal form.

6. Conclusion. We have studied the dynamics of a model PDE that is used in plasma physics to compare the merits of different discretization schemes for the speed variable in numerical solvers. To do so, we used the spectral transform associated with a singular Sturm–Liouville operator L to represent the solution in such a way that the dynamics becomes trivial, through multiplication by $e^{-\lambda t}$. Our algorithm relies on expressing the spectral density function of L in terms of the Titchmarsh–Weyl m -function, evaluating the m -function along a line segment in the complex plane, and extrapolating it to the real axis using Chebyshev polynomials. Our method is very general and will work for any singular Sturm–Liouville problem that is of limit point type at infinity and of limit circle (or regular) type at the origin, and if the m -function has enough smoothness near the real λ -axis to be well approximated by polynomials along line segments in the transverse direction. Furthermore, the

complexity of the new method for computing the m -function has been analyzed and shown to be comparable to other methods in double-precision, and faster in quadruple-precision, when optimal timestepping algorithms are used for all methods.

The solution of the PDE computed in this way can be compared with approximate solutions obtained with the discretization methods traditionally used in plasma physics and provides a basis for error quantification. Our construction of the solution and its behavior for certain initial conditions are also very helpful in explaining the behavior of approximate solutions. For example, we found that for singular initial conditions the solution of the PDE often cannot be resolved to the desired level of accuracy until t surpasses a critical value t^* , because the decay rate of the spectral transform of the solution is only algebraic at $t = 0$ and slow (but at least exponential) in the moments that follow. This has strong implications for the projected dynamics of this equation in finite-dimensional spaces of orthogonal polynomials, and for the choice of these polynomials, as we present elsewhere in subsequent work.

Appendix A. Bound on the number of steps. A detailed analysis of arbitrary-order ν -stage Runge–Kutta collocation methods of order 2ν , presented elsewhere [48], reveals that if such a scheme is used to evolve the fundamental matrix Φ_1 for

$$(A.1) \quad \begin{pmatrix} y \\ z \end{pmatrix}' = A(x) \begin{pmatrix} y \\ z \end{pmatrix}, \quad A(x) = \begin{pmatrix} 0 & \sqrt{\lambda/\Psi(x)} \\ (V(x) - \lambda)/\sqrt{\lambda\Psi(x)} & -\Psi'(x)/(2\Psi(x)) \end{pmatrix}$$

from x to $x + h$, with $\Phi_1(x) = I$ and $\lambda \in \mathbb{C}$, the local truncation error is bounded by

$$(A.2) \quad \|\Phi_{1,\text{numerical}}(x+h) - \Phi_{1,\text{exact}}(x+h)\| \leq (1.15\alpha^{-1})^{2\nu-1} e^{2rM(x)},$$

where $M(x) = 5|\lambda|^{1/2}\langle x \rangle^{3/2} + 3.5|\lambda|^{-1/2}\langle x \rangle^{1/2}$ is a bound on $\max_{|z-x| \leq 1/5} \|A(z)\|$ ($z \in \mathbb{C}, x \in \mathbb{R}$), $\langle x \rangle = \sqrt{1+x^2}$, α controls the size of a Bernstein ellipse [45] with foci at x and $x+h$, and semimajor and semiminor axes of length $h(\alpha + \alpha^{-1})/4$ and $h(\alpha - \alpha^{-1})/4$, respectively, and r is the radius of a disk centered at x in the complex plane containing the Bernstein ellipse, which requires $h/r \leq 4\alpha/(\alpha+1)^2$. To obtain (A.2), it is assumed in [48] that $\nu \geq 5$, $\alpha \geq 3$, $r \leq 1/5$, $hM \leq 2/3$. The first-order system (A.1) is equivalent to (2.12) and (2.16) when z is defined by $z(x) = \sqrt{\Psi(x)/\lambda}y'(x)$, which is scaled so that z grows at the same rate as y (like $x^{3/4}$) as $x \rightarrow \infty$ when λ is real. In practice, all three systems (2.12), (2.16), and (A.1) perform similarly, but the error analysis is simplest for (A.1).

We interpret (A.2) as a relative error in advancing any fundamental matrix $\Phi(x)$ from x to $x+h$ since $\Phi(x+h) = \Phi_1(x+h)\Phi(x)$. In floating point arithmetic with roundoff threshold δ , the accuracy of the result will cease to improve when this relative error reaches $O(\delta)$. The right-hand side of (A.2) will be less than δ if we assume $\nu \geq 5$, choose r so $rM = 5/3$, and require $\alpha \geq (5/3)(1/\delta)^{1/(2\nu-1)}$. Increasing α to 3 if necessary, the condition $h = 9r/(4\alpha)$ implies $h/r \leq 4\alpha/(\alpha+1)^2$, and we are led to the stepsize constraint

$$(A.3) \quad h(x)M(x) \leq \frac{1}{K_\nu(\delta)}, \quad K_\nu(\delta) = \max\left(\frac{3}{2}, \frac{4}{9}(1/\delta)^{1/(2\nu-1)}\right).$$

The intermediate assumptions that $r \leq 1/5$, $hM \leq 2/3$, and $\alpha \geq 3$ are ensured by (A.3) since $rM = 5/3$, $M(x) \geq 2\sqrt{5} \times 3.5\langle x \rangle$, and $\alpha = 9rM/(4hM) \geq 45/8$.

To derive (3.18), we note that counting steps using the largest stepsize allowed by (A.3) will give the points $x_{n+1} = x_n + h(x_n)$. This is Euler's method for the

ODE $dx/ds = h(x)$ with steps of size $\Delta s = 1$. The change in s after N steps is then N . Since $h(x) = [K_\nu M(x)]^{-1}$ is a positive, decreasing function, solutions of the continuous problem are increasing and concave down. Hence, Euler’s method will overpredict the solution $x(s)$ of this ODE. As a result, solving the ODE will overpredict the change in s needed to achieve a specified change in x using Euler’s method. Using separation of variables, we conclude that $\int_{x_1}^{x_2} K_\nu M(x) dx$ is an upper bound on the number of steps required to advance the solution of (A.1) from x_1 to x_2 with the maximum stepsize allowed by (A.3).

Appendix B. Technical lemmas. In this section we present four technical lemmas needed in Appendix C to establish the asymptotic behavior of the solutions u and y of (2.7) and (2.16), respectively.

LEMMA B.1. *The function $\Psi(x)$ in (2.2) is real analytic, even, and positive and satisfies $\Psi(x) \doteq 1/(2x^3)$, where $f(x) \doteq g(x)$ means that $f^{(n)}(x) - g^{(n)}(x) = o(x^{-k})$ as $x \rightarrow \infty$ for all integers $n \geq 0, k \geq 0$.*

Proof. Taylor expansion shows that $\Psi(z) = (2z^3)^{-1} [\operatorname{erf}(z) - (2/\sqrt{\pi})ze^{-z^2}]$ has a removable singularity at $z = 0$ with limiting value $\Psi(0) = 2/(3\sqrt{\pi}) \approx 0.3761$; thus, $\Psi(z)$ is entire. It is even since $2z^3, \operatorname{erf}(z)$, and ze^{-z^2} are odd. The formula

$$\Psi'(x) = -4\pi^{-1/2} \int_0^x (s/x)^4 e^{-s^2} ds < 0 \quad (x > 0)$$

shows that $\Psi(x)$ is decreasing on $(0, \infty)$. Since $\lim_{x \rightarrow \infty} \Psi(x) = 0$, it follows that $\Psi(x) > 0$ for $x \geq 0$. Since $\Psi(x)$ is even, it is positive for $x < 0$ as well. A straightforward induction argument shows that $g(x) = (2x^3)^{-1} - \Psi(x)$ has derivatives of the form

$$(B.1) \quad g^{(n)}(x) = \frac{(-1)^n}{\sqrt{\pi}} \left[\frac{(n+2)!}{2x^{n+3}} \int_x^\infty e^{-s^2} dx + \sum_{j=0}^n c_{nj} x^{2j-n-2} e^{-x^2} \right],$$

where $c_{00} = 1, c_{n0} = (n+1)(n!/2 + c_{n-1,0}), c_{nm} = 2c_{n-1,n-1}$, and

$$c_{nj} = (n+1-2j)c_{n-1,j} + 2c_{n-1,j-1} \quad (1 \leq j \leq n-1).$$

Since $\int_x^\infty e^{-s^2} dx \leq (2x)^{-1}e^{-x^2}$, there is a polynomial $p_n(x)$ of degree $n+1$ such that $|g^{(n)}(x)|x^k \leq p_n(x^2)x^{k-n-4}e^{-x^2}$, which converges to 0 as $x \rightarrow \infty$, as claimed. \square

LEMMA B.2. *Let $q \in \mathbb{R}$ and define $h(z) = z^q$. Suppose $g(x)$ and its derivatives grow slowly as $x \rightarrow \infty$; i.e., there exist integers $k_n \geq 0$ such that $g^{(n)}(x) = O(x^{k_n})$ for $n \geq 0$. Suppose also (increasing k_0 if necessary) that $g(x)^{-1} = O(x^{k_0})$. Then $f \doteq g$ implies $h \circ f \doteq h \circ g$, provided one of the following is true: q is an integer; g and f are real-valued; or the inverse of the distance from $g(x)$ to the negative real axis is $O(x^{k_0})$.*

Proof. By hypothesis, there exist $c_0 \geq 1, x_0 \geq 1$ such that

$$(B.2) \quad c_0^{-1}x^{-k_0} \leq |g(x)| \leq c_0x^{k_0} \quad (x \geq x_0).$$

Increasing x_0 if necessary, we may assume $|g(x) - f(x)| \leq (2c_0)^{-1}x^{-k_0} \leq c_0x^{k_0}$ for $x \geq x_0$. Any point ζ on the line segment $\gamma(x)$ joining $g(x)$ to $f(x)$ in the complex plane satisfies $|g(x) - \zeta| \leq |g(x) - f(x)|$, and hence

$$(B.3) \quad (2c_0)^{-1}x^{-k_0} \leq |\zeta| \leq (2c_0)x^{k_0} \quad (x \geq x_0, \zeta \in \gamma(x)).$$

In the complex case, if q is not an integer, we also have $|\operatorname{Im} g(x)| \geq c_0^{-1} x^{-k_0}$ whenever $x \geq x_0$ and $\operatorname{Re} g(x) \leq 0$; thus, $g(x)$ is closer to each $\zeta \in \gamma(x)$ than to the negative real axis and $\gamma(x)$ does not cross the branch cut of $h(z)$. If q is an integer, there is no branch cut. If f and g are real-valued, then $f(x)$ and $g(x)$ have the same sign for $x > x_0$ (since $g(x)$ is closer to $f(x)$ than to the origin), so $\gamma(x)$ is either a subset of the positive real axis or lies along the ray from the origin through $(-1)^q$. Either way, $\gamma(x)$ does not cross the branch cut. Next, for any $n \geq 0$ and $x \geq x_0$, we have

$$\left| h^{(n)}(f(x)) - h^{(n)}(g(x)) \right| = \left| \int_{\gamma(x)} h^{(n+1)}(\zeta) d\zeta \right| \leq \max_{\zeta \in \gamma(x)} |h^{(n+1)}(\zeta)| |f(x) - g(x)|.$$

Multiplying by x^k , where $k \geq 0$, and using (B.3), we obtain

$$|\dots| x^k \leq q(q-1)\dots(q-n)(2c_0 x^{k_0})^{|q-n-1|} x^k |f(x) - g(x)| \rightarrow 0, \quad (x \rightarrow \infty).$$

Since $h^{(n)}(g(x))$ grows slowly (i.e., polynomially) in x , the $k = 0$ case also implies that $h^{(n)}(f(x))$ grows slowly. Finally, we use Faà-di Bruno’s formula [27]

$$\frac{d^n}{dx^n} h(f(x)) = \sum_{\pi \in P_n} h^{(|\pi|)}(f(x)) \prod_{b \in \pi} f^{(|b|)}(x),$$

where P_n is the set of partitions of $\{1, 2, \dots, n\}$, $|\pi|$ is the number of blocks in the partition π , and $|b|$ is the number of integers in block b . Thus, for any integer $k \geq 0$,

$$\begin{aligned} & |(h \circ f)^{(n)}(x) - (h \circ g)^{(n)}(x)| x^k \\ \text{(B.4)} \quad & \leq \sum_{\pi \in P_n} \left| h^{(|\pi|)}(f(x)) \prod_{b \in \pi} f^{(|b|)}(x) - h^{(|\pi|)}(g(x)) \prod_{b \in \pi} g^{(|b|)}(x) \right| x^k. \end{aligned}$$

Subtracting and adding telescoping terms, e.g.,

$$|Aa_1a_2 - Bb_1b_2| \leq |Aa_1(a_2 - b_2)| + |A(a_1 - b_1)b_2| + |(A - B)b_1b_2|,$$

the right-hand side of (B.4) is bounded by a finite sum of terms in which one factor is a difference, either $|f^{(|b|)}(x) - g^{(|b|)}(x)|$ or $|h^{(|\pi|)}(f(x)) - h^{(|\pi|)}(g(x))|$, and the rest grow slowly in x . Since the difference converges to zero faster than any polynomial, the right-hand side of (B.4) converges to zero as $x \rightarrow \infty$, as claimed. \square

LEMMA B.3. *If $h(x)$ and its derivatives grow slowly, i.e., there exist integers $k_n \geq 0$ such that $h^{(n)}(x) = O(x^{k_n})$ for $n \geq 0$, then $f \doteq g$ implies $hf \doteq hg$. It suffices to check that $h \doteq H$ with $H^{(n)}(x) = O(x^{k_n})$.*

Proof. We see that $|([hf]^{(n)}(x) - [hg]^{(n)}(x)) x^k|$ may be bounded by

$$\text{(B.5)} \quad \sum_{j=0}^n \binom{n}{j} C_j x^{k_j+k} |f^{(n-j)}(x) - g^{(n-j)}(x)|$$

for large x , where C_j is a bound on $|h^{(j)}(x)| x^{-k_j}$ for large x . But (B.5) converges to zero as $x \rightarrow \infty$ due to $f^{(n-j)}(x) = g^{(n-j)}(x) + o(x^{k_j+k})$. Finally, we note that if $h \doteq H$, then $|h^{(n)}(x) - H^{(n)}(x)|$ can be made smaller than any multiple of x^{k_n} for large x , so if one is $O(x^{k_n})$, so is the other. \square

LEMMA B.4. *The solution $x(s)$ of $x' = \sqrt{\Psi(x)}$, $x(0) = 0$ exists for all $s \in \mathbb{R}$ and is an increasing, real analytic, odd function of s . There is a constant c such that $x(s) \doteq \sqrt[5]{25/8}(s - c)^{2/5}$.*

Proof. By Lemma B.1, $\Psi(x)$ is real analytic, even, and positive. It follows that $x(s)$ is increasing, odd, and real analytic for as long as the solution exists (see [10] regarding analyticity). Moreover, $s = \int_0^x \Psi(r)^{-1/2} dr$, which gives

$$(B.6) \quad s - \frac{2\sqrt{2}}{5}x^{5/2} = c - \int_x^\infty [\Psi(r)^{-1/2} - \sqrt{2}r^{3/2}] dr,$$

where $c = \int_0^\infty [\Psi(x)^{-1/2} - \sqrt{2}x^{3/2}] dx \approx 1.6247$. Since the integrand in (B.6) and each of its derivatives are $o(r^{-k})$ for all $k \geq 0$, we may apply Lemma B.2 to conclude

$$(B.7) \quad x = \sqrt[5]{25/8}(s - c)^{2/5} + \varepsilon_0(x), \quad \varepsilon_0(x) \doteq 0.$$

Since $dx/ds = \sqrt{\Psi(x)}$, we see that

$$\frac{d^n x}{ds^n} = \frac{d^n}{ds^n} \left[\sqrt[5]{25/8}(s - c)^{2/5} \right] + \varepsilon_n(x), \quad \varepsilon_n(x) = \sqrt{\Psi(x)} \frac{d}{dx} \varepsilon_{n-1}(x) \quad (n \geq 1).$$

By Lemmas B.1 and B.2, $\sqrt{\Psi(x)} \doteq (2x^3)^{-1/2}$, which has derivatives that grow slowly (in fact, decay). Thus, we may apply Lemma B.3 inductively to conclude that $\varepsilon_n(x) \doteq 0$ for $n \geq 0$. Since $\varepsilon_n(x) = o(x^{-k})$ for any $k \geq 0$, it follows from (B.7) that $\varepsilon_n(x(s)) = o(s^{-k})$ for $k \geq 0$, as claimed. \square

Appendix C. Asymptotics of the ODE. In this section, we study the asymptotic behavior of solutions of

$$(C.1) \quad -(\Psi w u')' = \lambda w u$$

for large x , where $\Psi(x) = [\operatorname{erf}(x) - x \operatorname{erf}'(x)]/(2x^3)$ and $w(x) = x^2 e^{-x^2}$. The case $\lambda = 0$ does not require asymptotic arguments as the general solution (2.8) can be written in closed form. We will show that if $\lambda \neq 0$, two linearly independent solutions of (C.1) exist of the form

$$(C.2) \quad u_\pm(x) = x^{-1/4} e^{x^2/2} P_0(x) \exp \left\{ \pm i \sqrt{\frac{8\lambda}{25}} P_1(x) x^{5/2} \right\} [1 + O(x^{-7/2})]$$

for $x \gg 1$, where

$$(C.3) \quad \begin{aligned} P_0(x) &= 1 + \frac{1}{8x\lambda} + \frac{5}{128x^2\lambda^2} + \frac{15}{1024x^3\lambda^3}, \\ P_1(x) &= 1 - \frac{5}{12x\lambda} - \frac{5}{32x^2\lambda^2} + \frac{5}{128x^3\lambda^3} + \frac{25}{6144x^4\lambda^4} + \frac{7 - 1152\lambda^4}{8192x^5\lambda^5}. \end{aligned}$$

When λ is real and positive, any real-valued solution of (C.1) may still be written as a linear combination $u(x) = Au_+(x) + Bu_-(x)$. Reality requires $B = \overline{A}$, which implies

$$(C.4) \quad u(x) = Cx^{-1/4} e^{x^2/2} P_0(x) \cos \left\{ \sqrt{8\lambda/25} P_1(x) x^{5/2} - \theta \right\} [1 + O(x^{-7/2})]$$

for some $C, \theta \in \mathbb{R}$, which yields (2.28). For all other values of $\lambda \in \mathbb{C} \setminus \{0\}$, one of the modes $u_\pm(x)$ in (C.2) grows superexponentially as $x \rightarrow \infty$ while the other decays. The recessive (decaying) mode is uniquely determined by (C.2), but the dominant (growing) mode is not. When $\lambda > 0$, neither mode dominates the other, so both are determined uniquely by (C.2).

While it is possible to derive (C.2) by guessing its form and computing successive terms of $P_0(x)$ and $P_1(x)$ from (C.1) iteratively, it is difficult to prove error bounds for the resulting series. Instead, we will use a Liouville transformation [7] to convert the ODE to normal form and use WKB theory [5, 33] to study the asymptotics.

To convert the general second-order self-adjoint equation $-(pu')' + qu = \lambda wu$ to Liouville normal form [7], $-U''(s) + Q(s)U(s) = \lambda U(s)$, one solves the ODE $dx/ds = \sqrt{p(x)/w(x)}$ to obtain $x(s)$ and then defines

$$U(s) = \gamma(s)u(x(s)), \quad \gamma(s) = \sqrt[4]{p(x(s))w(x(s))}, \quad Q(s) = \frac{q(x(s))}{w(x(s))} + \frac{\gamma''(s)}{\gamma(s)}.$$

To fit within the framework of WKB theory, it is convenient to change the sign of Q and absorb λ into the potential. Thus, we convert (C.1) to the form

$$(C.5) \quad -U'' = Q(s)U$$

by the change of variables $dx/ds = \sqrt{\Psi(x)}$, $U(s) = \Psi(x(s))^{1/4}y(x(s))$, $y(x) = w(x)^{1/2}u(x)$, $Q(s) = F(x(s))$, and

$$(C.6) \quad F(x) = \lambda - \left(\frac{1}{4}\Psi''(x) - \frac{1}{16}\Psi(x)^{-1}\Psi'(x)^2 + (1-x^2)\frac{\Psi'(x)}{x} + (x^2-3)\Psi(x) \right).$$

The same result is obtained if we start from (2.16) instead of (C.1). Since Ψ is even and entire, $\Psi'(z)/z$ has a removable singularity at $z = 0$. By Lemma B.1, $0 < \Psi(x) \leq 2/(3\sqrt{\pi})$ on the real axis. Thus, $F(x) - \lambda$ and $\sqrt{\Psi(x)}$ are real analytic on all of \mathbb{R} . By Lemma B.4, $Q(s) - \lambda$ is also real analytic. (We subtract λ to make $F(x)$ and $Q(s)$ real-valued for real arguments.) By Lemmas B.1, B.2, and B.3, we may replace $\Psi(x)$ by $1/(2x^3)$ in (C.6) to conclude

$$(C.7) \quad F(x) \doteq \lambda - \frac{1}{2}x^{-1} + \frac{9}{32}x^{-5},$$

where $f(x) \doteq g(x)$ means that $f^{(n)}(x) - g^{(n)}(x) = o(x^{-k})$ as $s \rightarrow \infty$ for all integers $n \geq 0$, $k \geq 0$. Lemmas B.2 and B.3 were both used (the latter twice) to convert $\Psi(x)^{-1}\Psi'(x)^2$ into $(9/2)x^{-5}$.

The WKB approximation [5, 33] of the solution of (C.5) is

$$(C.8) \quad U_{\pm}(s) \sim Q(s)^{-1/4} \exp \left\{ \pm i \int^s \sqrt{Q(r)} dr \right\},$$

where $f(s) \sim g(s)$ means $f(s)/g(s) \rightarrow 1$ as $s \rightarrow \infty$. When $Q(s)$ is real and positive, (C.8) can be derived by performing another Liouville transformation, namely, $d\xi/ds = \sqrt{Q(s)}$, $W(\xi(s)) = Q(s)^{1/4}U(s)$, to convert the ODE to $d^2W/d\xi^2 = [-1 + \phi(\xi)]W$. Neglecting ϕ gives $W \sim e^{\pm i\xi}$. If $\phi(\xi)$ is small, error estimates can be derived in the ξ coordinate system [33].

In our case, $Q(s)$ is complex-valued, so the change of variables $\xi = \int^s \sqrt{Q(r)} dr$ requires that we complexify the dependent variable [33]. We prefer to work with complex functions of the real variable x and only use s and ξ as intermediate steps to finding a representation of (C.1) that is suitable for perturbation analysis. To this end, we still define $W = Q^{1/4}U$ but treat it as a function of s rather than ξ . Substitution of the identity

$$(C.9) \quad \frac{d^2U}{ds^2} = Q^{1/4} \frac{d}{ds} \left[Q^{-1/2} \frac{d}{ds} (Q^{1/4}U) \right] + Q^{1/4} \left(\frac{d^2}{ds^2} Q^{-1/4} \right) U$$

into the equation $d^2U/ds^2 = -QU$ gives

$$(C.10) \quad Q^{-1/2} \frac{d}{ds} Q^{-1/2} \frac{dW}{ds} = [-1 + \phi(s)]W, \quad \phi(s) = -Q^{-3/4} \frac{d^2}{ds^2} Q^{-1/4},$$

where we have adopted the convention that differential operators act on all products that follow them unless otherwise indicated by parentheses. The left-hand side of (C.10) plays the role of $d^2W/d\xi^2$ in [33]. Changing back to the x -coordinate system and writing $W(s) = v(x(s))$, we find that if $v(x)$ and $u(x)$ are related by

$$(C.11) \quad u(x) = w(x)^{-1/2} [F(x)\Psi(x)]^{-1/4} v(x),$$

then $u(x)$ satisfies (C.1) iff $v(x)$ satisfies

$$(C.12) \quad \sqrt{\frac{\Psi}{F}} \frac{d}{dx} \sqrt{\frac{\Psi}{F}} \frac{dv}{dx} = \left(-1 + \sqrt{\frac{\Psi}{F}} \psi \right) v, \quad \psi = -F^{-1/4} \frac{d}{dx} \sqrt{\frac{\Psi}{F}} \frac{d}{dx} F^{-1/4}.$$

We again recognize the left-hand side as $d^2W/d\xi^2$.

Two technical issues concern zeros of $F(x)$ and the branch cut of the square root and fourth root functions along the negative real axis. We claim there is an $R_1 \geq 0$ such that $F(x)$ is bounded away from zero and neither $F(x)$ nor $\Psi(x)/F(x)$ crosses the negative real axis for $x \geq R_1$. Since we are interested in the asymptotics of $u(x)$ for large x , we need only solve (C.12) for $x \geq R_1$. If $\lambda \notin \mathbb{R}$, then $|F(x)| \geq |\text{Im } \lambda|$ since $\text{Im}\{F(x)\} = \text{Im } \lambda$, and $R_1 = 0$ suffices. If λ is real (and nonzero), then $F(x)$ may have zeros, but by (C.7) there is an $R_1 \geq 0$ such that for $x \geq R_1$, $F(x)$ and λ have the same sign, and $|F(x)| \geq |\lambda/2|$. If the sign is positive, the branch cut is avoided, and if the sign is negative, we treat $\text{Im } \lambda = \text{Im}\{F(x)\} = 0^+$ in all formulas involving fractional powers of λ or $F(x)$, e.g., $\sqrt{\Psi(x)/F(x)} = -i\sqrt{\Psi(x)/|F(x)|}$.

If ψ were zero in (C.12), $v(x) = e^{\pm i\xi(x)}$ would be independent solutions, where

$$(C.13) \quad \xi(x) = \int^x \sqrt{\frac{F(r)}{\Psi(r)}} dr.$$

Lemmas B.2 and B.3 justify replacing Ψ^{-1} by $2r^3$ and F by (C.7) to obtain an asymptotic formula for $\sqrt{F/\Psi}$:

$$(C.14) \quad \sqrt{\frac{F(r)}{\Psi(r)}} \doteq \sqrt{2\lambda} \left(1 - \frac{1}{2r\lambda} + \frac{9}{32r^5\lambda} \right)^{1/2} r^{3/2}.$$

We choose the integration constant in (C.13) so that

$$(C.15) \quad \xi(x) = (8\lambda/25)^{1/2} P_1(x)x^{5/2} + O(x^{-7/2}) \quad (x \rightarrow \infty),$$

where $P_1(x)$ was defined in (C.3). This result is obtained by expanding the asymptotic formula for $\sqrt{F/\Psi}$ in a binomial series and integrating term by term. We now look for solutions of (C.12) that are perturbations of $e^{\pm i\xi(x)}$:

$$(C.16) \quad v_+(x) = e^{i\xi(x)}(1 + h_+(x)), \quad v_-(x) = e^{-i\xi(x)}(1 + h_-(x)).$$

The functions $h_+(x)$ and $h_-(x)$ must satisfy

$$(C.17) \quad \frac{d}{dx} \sqrt{\frac{\Psi}{F}} \frac{dh}{dx} \pm 2i \frac{dh}{dx} = \psi(x)(1 + h(x)).$$

Solutions of the homogeneous problem (with $\psi \equiv 0$) are $h_{\pm} \equiv 1$ and $h_{\pm}(x) = e^{\mp 2i\xi(x)}$. Using variation of parameters [10] to solve the nonhomogeneous problem yields an integral equation for the solution with initial conditions $h(x_0) = h_0, h'(x_0) = 0$:

$$(C.18) \quad h_{\pm}(x) = h_0 \pm \frac{1}{2i} \int_{x_0}^x \left(1 - e^{\pm 2i\{\xi(y) - \xi(x)\}}\right) \psi(y)[1 + h_{\pm}(y)] dy.$$

The location of x_0 will be chosen below and depends on λ and the \pm sign. Simplifying the formula for ψ in (C.12) gives its behavior to leading order as $x \rightarrow \infty$:

$$(C.19) \quad \psi(x) = \frac{2\Psi'FF' - 5\Psi(F')^2 + 4\Psi FF''}{16\Psi^{1/2}F^{5/2}} \sim \frac{-7\sqrt{2}}{32\lambda^{3/2}x^{9/2}} \quad (x \rightarrow \infty).$$

Our goal is to use this in (C.18) to produce solutions $h_{\pm}(x)$ of (C.17) that decay like $O(x^{-7/2})$. If we are successful, then (C.2) will follow from (C.16), (C.11), (C.15), and

$$(C.20) \quad w(x)^{-1/2}[F(x)\Psi(x)]^{-1/4} = \frac{x^{-1/4}e^{x^2/2}P_0(x)[1 + O(x^{-4})]}{(\lambda/2)^{1/4}} \quad (x \rightarrow \infty),$$

with $P_0(x)$ as in (C.3). Indeed, the factor of $(\lambda/2)^{-1/4}$ can be dropped by linearity, and the factors of $1 + O(x^{-4}), \exp\{O(x^{-7/2})\}$, and $1 + h_{\pm}(x)$ combine to make $1 + O(x^{-7/2})$ in (C.2), as claimed. The coefficients of $P_0(x)$ were obtained from (C.20) using (C.7), $\Psi(x) \doteq (2x^3)^{-1}$, Lemmas B.2 and B.3, and the binomial series.

Let us therefore study solutions of the integral equation (C.18). Let $\alpha = \text{Im} \sqrt{\lambda}$ and $\beta = \pm 1$, depending on the case considered in (C.16). We drop the \pm subscript on h and define

$$K(x, y) = \frac{\beta}{2i} [1 - E(x, y)]\psi(y), \quad E(x, y) = e^{2i\beta\{\xi(y) - \xi(x)\}} \quad (\beta = \pm 1),$$

which appear in (C.18). We claim that there is an $R \geq R_1$ (depending on λ and β) such that one of the following holds:

- case 1: $y \geq x \geq R \Rightarrow |K(x, y)| \leq |\psi(y)|,$
- case 2: $|\alpha| > 0$ and $(x \geq y \geq R \Rightarrow |K(x, y)| \leq |\psi(y)|, |E(x, y)| \leq \eta(x, y)),$

where $\eta(x, y) = \exp\{-2|\alpha|y^{3/2}(x - y)\}$. If $\alpha = 0$, then $\lambda > 0$, and setting $R = R_1$ suffices to establish case 1. Indeed, since $F(x) > 0$ for $x \geq R_1$, it follows that $\xi(x)$ in (C.13) is real, $E(x, y)$ is on the unit circle, and $|1 - E| \leq 2$ for $x \geq R, y \geq R$. If $\alpha > 0$, we see from (C.14) that $\text{Im} \sqrt{F(r)/\Psi(r)} r^{-3/2} \rightarrow (\sqrt{2}\alpha)$ as $r \rightarrow \infty$. Since $\sqrt{2} > 1$, there exists $R \geq R_1$ such that

$$\text{Im} \sqrt{F(r)/\Psi(r)} \geq \alpha r^{3/2} \quad (r \geq R).$$

It then follows from (C.13) that

$$(C.21) \quad \text{Im}\{\xi(y) - \xi(x)\} \geq \int_x^y \alpha r^{3/2} dr \geq \alpha x^{3/2}(y - x) \geq 0 \quad (y \geq x \geq R).$$

Since $E(x, y) = e^{2i\beta\{\xi(y) - \xi(x)\}}$, we see that $E(x, y)$ lies inside or on the unit circle when $\beta = 1$ and $y \geq x \geq R$. Thus, case 1 holds when $\beta = 1$. Interchanging x and y in (C.21) and evaluating $|E(x, y)|$ shows that case 2 holds when $\beta = -1$. Similar arguments show that if $\alpha < 0$, cases 1 and 2 hold when $\beta = -1$ and $\beta = 1$, respectively.

In case 1, we set $h_0 = 0$ in (C.18) and send x_0 to infinity. This yields the equation

$$h(x) = \mathbb{K}[1 + h](x), \quad \mathbb{K}[f](x) = - \int_x^\infty K(x, y)f(y) dy.$$

Increasing R if necessary, we may assume $\int_R^\infty |\psi(y)| dy \leq 1/2$. Then $\|\mathbb{K}\| \leq 1/2$, where \mathbb{K} is regarded as an operator on $BC([R, \infty))$, the Banach space of bounded, continuous functions in the uniform norm. Thus, $h = (\mathbb{K} + \mathbb{K}^2 + \mathbb{K}^3 + \dots)1$ is the unique bounded, continuous function that satisfies the integral equation. By the dominated convergence theorem and Leibniz integral rule, such a solution of the integral equation also satisfies the ODE (C.17). We note that $h_n = (\mathbb{K} + \mathbb{K}^2 + \dots + \mathbb{K}^n)1$ can be computed via the Picard iteration $h_0(x) = 0, h_{n+1} = \mathbb{K}[1 + h_n]$. Standard estimates [10, 33] on the size of $|h_{n+1}(x) - h_n(x)|$ in the Picard iteration scheme give the bound

$$(C.22) \quad |h(x)| \leq \exp\left(\int_x^\infty |\psi(r)| dr\right) - 1 \quad (x > R).$$

By (C.19), $h(x) = O(x^{-7/2})$ as $x \rightarrow \infty$, as required.

In case 2, we define $\mathbb{K}[f](x) = \int_R^x K(x, y)f(y) dy$ and proceed in the same manner, again assuming $\int_R^\infty |\psi(y)| dy \leq 1/2$. This establishes existence and uniqueness of a bounded, continuous solution of $h = \mathbb{K}[1 + h]$, along with the bound

$$(C.23) \quad |h(x)| \leq \left(\exp\left\{\int_R^x |\psi(r)| dr\right\} - 1\right) \leq (e^{1/2} - 1) \leq 2/3 \quad (x \geq R).$$

Although $h(x)$ will not in general approach 0 as $x \rightarrow \infty$, we will show below that it approaches a limiting value, c , with $|c| \leq 2/3$. Defining

$$(C.24) \quad \tilde{h}(x) = \frac{h(x) - c}{1 + c}, \quad h_0 = \frac{-c}{1 + c},$$

we find that \tilde{h} satisfies $\tilde{h} = h_0 + \mathbb{K}[1 + \tilde{h}]$ and hence (C.17). Thus, adjusting the initial condition from 0 to h_0 merely shifts and rescales the solution. If we can show that $h(x) - c = O(x^{-7/2})$ as $x \rightarrow \infty$, then $\tilde{h}(x)$ will be the desired solution of (C.17) that decays as $O(x^{-7/2})$. To prove that $c = \lim_{x \rightarrow \infty} h(x)$ exists, we integrate (C.17) from x_1 to x_2 , assuming $x_2 > x_1 \geq R$:

$$(C.25) \quad h(x_2) - h(x_1) = \frac{\beta}{2i} \int_{x_1}^{x_2} \psi(x)(1 + h(x)) dx - \frac{\beta}{2i} \sqrt{\frac{\Psi(x)}{F(x)}} h'(x) \Big|_{x_1}^{x_2}.$$

Recall that $\beta = \pm 1$ distinguishes the case in (C.16). The first term on the right is $O(x_1^{-7/2})$ due to (C.23) and (C.19). If we can show that $\sqrt{\Psi/F}h'(x) = O(x^{-7/2})$, then we are done: (C.25) implies that for any sequence $x_1 < x_2 < \dots$ with $x_n \rightarrow \infty$, $h(x_n)$ is a Cauchy sequence. So $c = \lim_{x \rightarrow \infty} h(x)$ exists. Sending x_2 to ∞ and replacing x_1 by x in (C.25) then gives $c - h(x) = O(x^{-7/2})$, as required.

To show that $\sqrt{\Psi/F}h'(x) = O(x^{-7/2})$, we differentiate the integral equation $h = \mathbb{K}[1 + h]$. This gives

$$(C.26) \quad \sqrt{\frac{\Psi(x)}{F(x)}} h'(x) = \int_R^x E(x, y)\psi(y)[1 + h(y)] dy \quad (x > R).$$

Using the bound $|E(x, y)| \leq \eta(x, y) = \exp\{-2|\alpha|y^{3/2}(x - y)\}$ for $x \geq y \geq R$, and breaking the integral in (C.26) into two segments of length $(x - R)/2$, we obtain

$$(C.27) \quad \left| \sqrt{\frac{\Psi(x)}{F(x)}} h'(x) \right| \leq \frac{5}{3} \eta\left(x, \frac{x+R}{2}\right) \int_R^{\frac{x+R}{2}} |\psi(y)| dy + \frac{5}{3} \int_{\frac{x+R}{2}}^x |\psi(y)| dy,$$

where $5/3$ is a bound on $|1 + h(y)|$. The first integral on the right is bounded by $1/2$, while

$$\eta\left(x, \frac{x+R}{2}\right) = \exp\left\{-|\alpha| \left(\frac{x+R}{2}\right)^{1/2} \left(\frac{x^2 - R^2}{2}\right)\right\},$$

which decays superexponentially as $x \rightarrow \infty$. Since $\int_x^\infty |\psi(y)| dy = O(x^{-7/2})$, the second integral in (C.27) is bounded by

$$C \left(\frac{x+R}{2}\right)^{-7/2} = 2^{7/2} C \left(1 + \frac{R}{x}\right)^{-7/2} x^{-7/2} = O(x^{-7/2}).$$

Thus, $\sqrt{\Psi/F} h'(x) = O(x^{-7/2})$, as claimed.

Appendix D. Analyticity of the spectral density function. A number of authors have proved that the spectral function $\rho(x)$ is absolutely continuous when the potential in the Schrödinger equation is of bounded variation or decreases sufficiently rapidly at infinity; see, e.g., [35, 46, 30]. However, we are not aware of any work that establishes conditions to ensure that $\rho'(\lambda)$ will be real analytic for $\lambda > 0$. In this appendix, we prove this for the operator L in (2.4), studied throughout this paper.

Our proof will be to show that $u_\pm(x; \lambda)$ in (C.2) can be made to depend analytically on λ in a complex neighborhood of each $\lambda_0 > 0$. We can then construct a fundamental matrix $\tilde{\Phi}(x; \lambda)$ for (3.14) of the form

$$\tilde{\Phi}(x; \lambda) = \begin{pmatrix} y_+(x; \lambda) & y_-(x; \lambda) \\ \Psi(x)y'_+(x; \lambda) & \Psi(x)y'_-(x; \lambda) \end{pmatrix},$$

where $y_\pm(x; \lambda) = xe^{-x^2/2}u_\pm(x; \lambda)$. The construction of $u_\pm(x; \lambda)$ will involve a fixed-point (i.e., Picard) iteration, as in Appendix C, producing solutions for sufficiently large x , say, $x \geq x_0$, where x_0 does not depend on λ in the neighborhood. Similar to what we did in (3.9), we can express the fundamental matrix $\Phi(x; \lambda)$ with correct boundary conditions at $x = 0$ in the form

$$(D.1) \quad \Phi(x; \lambda) = \tilde{\Phi}(x; \lambda)C(\lambda) \quad (x \geq x_0),$$

where $C(\lambda) = \tilde{\Phi}(x_0; \lambda)^{-1}\Phi(x_0; \lambda)$ depends analytically on λ near λ_0 , since $\Phi(x_0; \lambda)$ is an entire function of λ (see [10]) while $\tilde{\Phi}(x_0; \lambda)^{-1}$ is analytic in the neighborhood where $u_\pm(x; \lambda)$ are analytic. The m -function may then be written

$$m(\lambda) = - \lim_{x \rightarrow \infty} \frac{y_+(x; \lambda)c_{11}(\lambda) + y_-(x; \lambda)c_{21}(\lambda)}{y_+(x; \lambda)c_{12}(\lambda) + y_-(x; \lambda)c_{22}(\lambda)} = \begin{cases} -c_{21}(\lambda)/c_{22}(\lambda), & \text{Im } \lambda > 0, \\ -c_{11}(\lambda)/c_{12}(\lambda), & \text{Im } \lambda < 0. \end{cases}$$

This equation clarifies the source of the discontinuity across the continuous spectrum in the m -function: when $\text{Im } \lambda$ changes sign, the growing and decaying solutions switch,

i.e., $y_+(x; \lambda)$ and $y_-(x; \lambda)$ reverse roles. From (2.22), we know that $\overline{m(\lambda)} = m(\bar{\lambda})$. Thus,

$$(D.2) \quad \rho'(\lambda) = \frac{1}{\pi} \operatorname{Im}\{m(\lambda^+)\} = \frac{1}{\pi} \frac{m(\lambda^+) - m(\lambda^-)}{2i} = \frac{1}{2\pi i} \frac{\det C(\lambda)}{c_{12}(\lambda)c_{22}(\lambda)},$$

which is real-valued for real λ and analytic in the neighborhood on which $u_{\pm}(x; \lambda)$ can be constructed to depend analytically on λ . Reality of the final formula in (D.2) for $\lambda \in \mathbb{R}$ can be confirmed by noting that $C(\lambda) = \check{\Phi}^{-1}\Phi$ with Φ real and $\check{\Phi}$ having complex conjugate columns (since $u_+(x; \lambda) = u_-(x; \bar{\lambda})$). We note that (D.2) could potentially be used to compute $\rho'(\lambda)$ without complexifying λ , but by solving a terminal value problem as well as an initial value problem. This is the key idea of the FPP algorithm [15], though their derivation is different than ours and leads to a representation, namely, (5.1), that does not allow λ to be complex.

The WKB analysis in Appendix C fails to produce analytic functions $u_{\pm}(x; \lambda)$ because different formulas are used for the cases when $\operatorname{Im} \lambda$ is positive or negative. The problem occurs in case 2, where we adjusted h_0 in (C.18) to achieve $\lim_{x \rightarrow \infty} h_{\pm}(x) = 0$. In fact, (C.18) is not the most general solution of (C.17); one could add to it a term of the form $\pm \frac{h_1}{2i} (1 - e^{\pm 2i\{\xi(x_0) - \xi(x)\}}) \sqrt{\Psi(x_0)/F(x_0)}$, so that $h_{\pm}(x_0) = h_0$ and $h'_{\pm}(x_0) = h_1$. Varying h_1 and solving for h_0 (as we did with $h_1 = 0$) leads to a one-parameter family of solutions of (C.17) that approach zero as $x \rightarrow \infty$, consistent with the observation at the beginning of Appendix C that the recessive solution is unique while the dominant solution is not. Our task now is to analytically continue $u_+(x; \lambda)$ from the upper λ -half-plane into the lower half-plane, and $u_-(x; \lambda)$ from the lower half-plane into the upper half-plane, in order to select the “right” dominant solutions.

It will be necessary to complexify x in Appendix C, as well as in Lemmas B.1–B.3. We will use the letter z (to replace x) as there is no chance of confusion with the second component of \vec{r} in (2.12). First, we modify the notation $f(z) \doteq g(z)$ to mean that there is a number $R \geq 0$ such that $f(z)$ and $g(z)$ are both defined and analytic on the region $S_R = \{re^{i\theta} : r > R, |\theta| < \pi/6\}$, and, for all nonnegative integers n and k , $f^{(n)}(z) - g^{(n)}(z) = o(|z|^{-k})$ as $(z \in S_R) \rightarrow \infty$, i.e., as $|z| \rightarrow \infty$ with z remaining in S_R . We then have the following lemma.

LEMMA D.1. $\Psi(z)$ satisfies $\Psi(z) \doteq 1/(2z^3)$.

Proof. Since $\Psi(z)$ is entire, both $\Psi(z)$ and $(2z)^{-3}$ are defined on S_R with $R = 0$. Proceeding as in Lemma B.1, we define $g(z) = (2z)^{-1} - \Psi(z)$ and observe that (B.1) holds with x replaced by $z = x + iy$ and the integral interpreted as $\int_x^\infty e^{-(s+iy)^2} ds$. Note that $|\int_x^\infty e^{-(s+iy)^2} ds| \leq e^{y^2} \int_x^\infty e^{-s^2} ds \leq (2x)^{-1} e^{y^2 - x^2}$. Since $z \in S$, we have $y^2 \leq x^2/3$, $|z|^2 \leq 4x^2/3$, and $x^2 - y^2 \geq \frac{2}{3}x^2 \geq \frac{1}{2}|z|^2$. It follows that $(2x)^{-1} e^{y^2 - x^2} \leq (\sqrt{3}|z|)^{-1} e^{-|z|^2/2}$. Similarly, $|e^{-z^2}| \leq e^{-|z|^2/2}$. The rest of the proof of Lemma B.1 works the same, i.e., there is a polynomial $p_n(x)$ of degree $n+1$ such that $|g^{(n)}(z)||z|^k \leq p_n(|z|^2)|z|^{k-n-4} e^{-|z|^2/2}$, which converges to zero as $(z \in S_R) \rightarrow \infty$. \square

Since $g(z) = 1/(2z^3)$ satisfies $\frac{1}{2}|z|^{-3} \leq |g(z)| \leq 2|z|^{-3}$, which is of the form (B.2), the proof of Lemma B.2 is easily modified to show $\Psi(z)^{-1} \doteq 2z^3$. In particular, there is an R large enough that $\Psi(z)^{-1}$ has no poles for $z \in S_R$, which also follows from $|\Psi(z)| \geq |2z^3|^{-1} - |\Psi(z) - (2z^3)^{-1}|$ and Lemma D.1. Adapting Lemma B.3 to the case of analytic functions on S_R is also straightforward and implies that $F(z; \lambda) = \lambda - [\frac{1}{4}\Psi''(z) - \frac{1}{16}\Psi(z)^{-1}\Psi'(z)^2 + (1 - z^2)z^{-1}\Psi'(z) + (z^2 - 3)\Psi(z)]$ satisfies

$$(D.3) \quad F(z; \lambda) - \lambda \doteq -\frac{1}{2}z^{-1} + \frac{9}{32}z^{-5},$$

where the left-hand side is independent of λ . Next we fix $\lambda_0 > 0$ and consider $\lambda \in B$, where $B = B_a(\lambda_0)$ is the open ball of radius $a = \lambda_0/3$ centered at λ_0 . By (D.3), we may increase R if necessary so that $|F(z; \lambda) - \lambda| < \lambda_0/6$. We then have $|F(z; \lambda) - \lambda_0| < \lambda_0/6 + \lambda_0/3 = \lambda_0/2$, which implies that $|\text{Arg } F(z; \lambda)| < \pi/6$ for $\lambda \in B$ and $z \in S_R$. Since $\Psi(z)^{-1} \doteq 2z^3$, we may increase R further if necessary to conclude that $\Psi(z)^{-1} = 2z^3(1 + \varepsilon(z))$ for $z \in S_R$, where $|\varepsilon(z)| < \sin(\pi/12)$. As a result, $|\text{Arg } \Psi(z)| = |\text{Arg } [\Psi(z)^{-1}]| < 7\pi/12$ for $z \in S_R$. Thus, $F(z; \lambda)$, $\Psi(z)^{-1}$ and $F(z; \lambda)/\Psi(z)$ are bounded away from the branch cut of the square root and fourth root functions for $\lambda \in B$ and $z \in S_R$. If we define

$$v(z; \lambda) = w(z)^{1/2} [F(z; \lambda)/\lambda]^{1/4} [2\Psi(z)]^{1/4} u(z; \lambda),$$

then $u(z; \lambda)$ will satisfy (C.1) iff $v(z; \lambda)$ satisfies

$$(D.4) \quad \sqrt{\frac{\Psi}{F}} \frac{d}{dz} \sqrt{\frac{\Psi}{F}} \frac{dv}{dz} = \left(-1 + \sqrt{\frac{\Psi}{F}} \psi \right) v, \quad \psi = -F^{-1/4} \frac{d}{dz} \sqrt{\Psi} \frac{d}{dz} F^{-1/4}.$$

This may be derived analogously to (C.11) and (C.12). We also define

$$(D.5) \quad \xi(z; \lambda) = \int^z \sqrt{\frac{F(\zeta; \lambda)}{\Psi(\zeta)}} d\zeta \quad (\lambda \in B, z \in S_R)$$

and look for solutions of (D.4) of the form $v_{\pm}(z; \lambda) = e^{\pm i\xi(z; \lambda)} [1 + h_{\pm}(z; \lambda)]$. The equation for $h_{\pm}(z; \lambda)$ is

$$(D.6) \quad \frac{d}{dz} \sqrt{\frac{\Psi}{F}} \frac{dh}{dz} \pm 2i \frac{dh}{dz} = \psi(z; \lambda)(1 + h), \quad h = h_{\pm}(z; \lambda).$$

We can solve this equation for $\lambda \in B$ and $z \in S_R^{\pm}$, where

$$(D.7) \quad S_R^+ = \{r e^{i\theta} : r > R, -\pi/24 < \theta < \pi/6\}, \quad S_R^- = \{\bar{z} : z \in S_R^+\}.$$

This is done by solving the integral equation

$$h(z; \lambda) = - \int_{\gamma(z)} K(z, \zeta; \lambda) [1 + h(\zeta; \lambda)] d\zeta,$$

where $K(z, \zeta; \lambda) = \frac{\beta}{2i} [1 - E(z, \zeta; \lambda)] \psi(\zeta; \lambda)$, $E(z, \zeta; \lambda) = e^{2i\beta\{\xi(\zeta; \lambda) - \xi(z; \lambda)\}}$, $\gamma(z)$ is the path from z to ∞ in the direction $e^{i\beta\pi/6}$, and $\beta = \pm 1$ records the sign in (D.6). Thus, we seek a solution of $h = \mathbb{K}[1 + h]$, where

$$(D.8) \quad \mathbb{K}f(z; \lambda) = - \int_0^{\infty} K(z, z + \hat{\gamma}s; \lambda) f(z + \hat{\gamma}s; \lambda) \hat{\gamma} ds, \quad \hat{\gamma} = e^{i\beta\pi/6}.$$

Focusing on the $\beta = +1$ case, we will show below (increasing R if necessary) that there is a positive, decreasing function $g(s)$ defined for $s \geq 0$ such that $\int_0^{\infty} g(s) ds \leq 1/2$ and $|K(z, z + \hat{\gamma}s; \lambda)| \leq g(s)$ for $z \in S_R^+$, $\lambda \in B$, and $s \geq 0$.

Let \mathcal{B} be the Banach space of (jointly) holomorphic functions $f(z; \lambda)$ that are bounded on $S_R^+ \times B$, with norm $\|f\|_{\mathcal{B}} = \sup_{z, \lambda} |f(z; \lambda)|$. Then since the integrand of (D.8) is holomorphic in z and λ for fixed s and is uniformly dominated by $g(s)\|f\|_{\mathcal{B}}$, \mathbb{K} maps \mathcal{B} to \mathcal{B} and has norm $\leq 1/2$. Indeed, continuity of $\mathbb{K}f(z; \lambda)$ follows from

the dominated convergence theorem, and analyticity follows from Morera’s theorem and Fubini’s theorem. The fixed-point iteration $h = [\mathbb{K} + \mathbb{K}^2 + \mathbb{K}^3 + \dots]1$ leads to a holomorphic function $h \in \mathcal{B}$ that satisfies $h = \mathbb{K}[1 + h]$ and $\|h\|_{\mathcal{B}} \leq 1$. A change of variables shows that for any $t \geq 0$ we have

$$h(z + \hat{\gamma}t; \lambda) = - \int_t^\infty K(z + \hat{\gamma}t, z + \hat{\gamma}s; \lambda)[1 + h(z + \hat{\gamma}s; \lambda)]\hat{\gamma} ds.$$

Applying $\hat{\gamma}^{-1}\partial_t[\dots]$ and $\hat{\gamma}^{-1}\partial_t\{\sqrt{\Psi/F}\hat{\gamma}^{-1}\partial_t[\dots]\}$ to this equation and setting $t = 0$ shows that $h(z; \lambda)$ satisfies (D.6). Finally, we will see below that

$$(D.9) \quad |K(z, z + \hat{\gamma}s; \lambda)| \leq g(s + 3(|z| - R)/4) \quad (z \in S_R^+, \lambda \in B).$$

It follows that $|h(z; \lambda)| \leq 2 \int_{3(|z|-R)/4}^\infty g(s) ds$, which converges to zero as $|z| \rightarrow \infty$. In particular, for real x , $h(x; \lambda) \rightarrow 0$ as $x \rightarrow \infty$. As a result, $v_+(x; \lambda) \sim e^{+i\xi(x; \lambda)}$ for large x , and $u_+(x; \lambda)$ has the form (C.2) and depends analytically on $\lambda \in B$ for fixed x . In the $\beta = -1$ case, the same construction works on S_R^- , and in fact the partial sums $h_-^{(n)} = \sum_{k=1}^n \mathbb{K}^k 1$ are related to those above by $h_-^{(n)}(z; \lambda) = \overline{h_+^{(n)}(\bar{z}; \bar{\lambda})}$ for $z \in S_R^-$, $\lambda \in B$. As a result, $u_-(x; \bar{\lambda}) = \overline{u_+(x; \lambda)}$ for $x > R$. In (D.1), the point x_0 where $C(\lambda) = \tilde{\Phi}(x_0; \lambda)^{-1}\Phi(x_0; \lambda)$ is defined can be any number greater than R .

It remains to construct $g(s)$ so that $\int_0^\infty g(s) ds \leq 1/2$ and (D.9) holds. First we claim that $|E(z, z + \hat{\gamma}s; \lambda)|$ is a decreasing function of s when $\lambda \in B$ and $z \in S_R^\pm$ are fixed and therefore remains bounded by 1 for $s \geq 0$. Focusing on the $\beta = +1$ case, this is equivalent to claiming that $\text{Im}\{\xi(z + \hat{\gamma}s; \lambda)\}$ is an increasing function of s . This follows from $(d/ds)\xi(z + \hat{\gamma}s; \lambda) = \sqrt{F(\zeta; \lambda)}/\Psi(\zeta)\hat{\gamma}$, where $\zeta = z + \hat{\gamma}s$. Indeed, since $\zeta \in S_R^+$ and $\lambda \in B$, the arguments used above to bound $F(z; \lambda)$ and $\Psi(z)^{-1}$ away from the branch cut also imply

$$\begin{aligned} \text{Arg} \left(\sqrt{\frac{F(\zeta; \lambda)}{\Psi(\zeta)}} \hat{\gamma} \right) &= \frac{1}{2} \text{Arg } F + \frac{3}{2} \text{Arg } z + \frac{1}{2} \text{Arg}(1 + \varepsilon(z)) + \text{Arg } \hat{\gamma} \\ &\in \left(-\frac{\pi}{12} - \frac{\pi}{16} - \frac{\pi}{24} + \frac{\pi}{6}, \frac{\pi}{12} + \frac{\pi}{4} + \frac{\pi}{24} + \frac{\pi}{6} \right) = \left(\frac{\pi}{48}, \frac{13\pi}{24} \right). \end{aligned}$$

Therefore, the imaginary part of $(d/ds)\xi(z + \hat{\gamma}s; \lambda)$ is positive. Note that $\text{Arg } \hat{\gamma} = \pi/6$ offsets the negative contributions from the other terms, which is why adjusting the contour of integration to point in the $\hat{\gamma}$ direction allows us to analytically continue λ across the real axis.

Since $|E(z, z + \hat{\gamma}s; \lambda)| \leq 1$ for $s \geq 0$, the left-hand side of (D.9) is bounded by $|\psi(z + \hat{\gamma}s; \lambda)|$. Increasing R if necessary, we claim there is a nonnegative function $g(s)$ such that $\int_0^\infty g(s) ds \leq 1/2$ and $|\psi(z + \hat{\gamma}s; \lambda)| \leq g(s + 3(|z| - R)/4) \leq g(s)$ for $s \geq 0$, $z \in S_R^+$, $\lambda \in B$. (We continue to assume $\beta = +1$.) To this end, we note that

$$\begin{aligned} \psi(\zeta; \lambda) &= \frac{1}{8} \underbrace{\Psi^{-1/2}\Psi'F'}_{C_1} F^{-3/2} - \frac{5}{16} \underbrace{\Psi^{1/2}(F')^2}_{C_2} F^{-5/2} + \frac{1}{4} \underbrace{\Psi^{1/2}F''}_{C_3} F^{-3/2}, \\ C_1 &\doteq -\frac{3}{\sqrt{8}}\zeta^{-9/2}[1 - \frac{45}{16}\zeta^{-4}], \quad C_2 \doteq \frac{\zeta^{-11/2}}{\sqrt{32}}[1 - \frac{45}{16}\zeta^{-4}]^2, \quad C_3 \doteq -\frac{\zeta^{-9/2}}{\sqrt{2}}[1 - \frac{135}{16}\zeta^{-4}], \end{aligned}$$

where we used (D.3) to compute F' and F'' , which are independent of λ . Increasing R if necessary, we may assume $|C_1| \leq (4/3)|\zeta|^{-9/2}$, $|C_2| \leq (1/5)|\zeta|^{-11/2}$, and $|C_3| \leq$

$(4/5)|\zeta|^{-9/2}$ for $\zeta \in S_R$. We already established that $|F(\zeta; \lambda) - \lambda_0| < \lambda_0/2$ for $\lambda \in B$ and $\zeta \in S_R$, so $|F(\zeta; \lambda)| > \lambda_0/2$. It follows that

$$(D.10) \quad |\psi(\zeta; \lambda)| \leq \frac{11}{30} \left(\frac{2}{\lambda_0}\right)^{3/2} |\zeta|^{-9/2} + \frac{1}{16} \left(\frac{2}{\lambda_0}\right)^{5/2} |\zeta|^{-11/2} \quad (\zeta \in S_R, \lambda \in B).$$

Finally, for $z \in S_R^+$ and $\zeta = z + \hat{\gamma}s$ with $s \geq 0$, we have $|\zeta|^2 = |z|^2 + s^2 - 2|z|s \cos[5\pi/6 + \text{Arg}(z)]$ by the law of cosines. Since $-\pi/24 < \text{Arg}(z) < \pi/6$ and $\cos(19\pi/24) < -3/4$, $|\zeta|^2 \geq |z|^2 + s^2 + (3/2)|z|s$. It follows that $|\zeta| \geq (s + 3|z|/4)$. Defining $g(s)$ to be the right-hand side of (D.10) with $|\zeta|$ replaced by $(s + 3R/4)$, we have that $g(s)$ is a positive, decreasing function for $s \geq 0$ and $|\psi(z + \hat{\gamma}s; \lambda)| \leq g(s + 3(|z| - R)/4) \leq g(s)$ for $z \in S_R^+$ and $\lambda \in B$, as claimed. Increasing R if necessary, $\int_0^\infty g(s) ds \leq 1/2$.

REFERENCES

- [1] I.G. ABEL, M. BARNES, S.C. COWLEY, W. DORLAND, AND A.A. SCHEKOCIHIN, *Linearized model Fokker–Planck collision operators for gyrokinetic simulations. I. Theory*, Phys. Plasmas, 15 (2008), 122509.
- [2] M. BARNES, I.G. ABEL, W. DORLAND, D.R. ERNST, G.W. HAMMETT, P. RICCI, B.N. ROGERS, A.A. SCHEKOCIHIN, AND T. TATSUNO, *Linearized model Fokker–Planck collision operators for gyrokinetic simulations. II. Numerical implementation and tests*, Phys. Plasmas, 16 (2009), 072107.
- [3] M. BARNES, I.G. ABEL, W. DORLAND, T. GÖRLEN, G.W. HAMMETT, AND F. JENKO, *Direct multiscale coupling of a transport code to gyrokinetic turbulence codes*, Phys. Plasmas, 17 (2010), 056109.
- [4] M. BARNES, W. DORLAND, AND T. TATSUNO, *Resolving velocity space dynamics in continuum gyrokinetics*, Phys. Plasmas, 17 (2010), 032106.
- [5] C.M. BENDER AND S.A. ORSZAG, *Advanced Mathematical Methods for Scientists and Engineers I, Asymptotic Methods and Perturbation Theory*, 2nd ed., Springer, New York, 1999.
- [6] C. BENNEWITZ AND W.N. EVERITT, *The Titchmarsh–Weyl eigenfunction expansion theorem for Sturm–Liouville differential equations*, in Sturm–Liouville Theory, Past and Present, W. O. Amrein, A. M. Hinz, and D. P. Pearson, eds., Birkhäuser, Basel, 2005.
- [7] G. BIRKHOFF AND G.-C. ROTA, *Ordinary Differential Equations*, Ginn–Blaisdell, Waltham, MA, 1969.
- [8] V. BRATANOV, F. JENKO, D. HATCH, AND S. BRUNNER, *Aspects of linear Landau damping in discretized systems*, Phys. Plasmas, 20 (2013), 022108.
- [9] J. CANDY, C. HOLLAND, R.E. WALTZ, M.R. FAHEY, AND E. BELLI, *Tokamak profile prediction using direct gyrokinetic and neoclassical simulation*, Phys. Plasmas, 16 (2009), 060704.
- [10] E.A. CODDINGTON AND N. LEVINSON, *Theory of Ordinary Differential Equations*, Krieger Publishing, Malabar, FL, 1984.
- [11] W.F. DONOGHUE, JR., *Monotone Matrix Functions and Analytic Continuation*, Springer-Verlag, New York, 1974.
- [12] A. DUTT, L. GREENGARD, AND V. ROKHLIN, *Spectral deferred correction methods for ordinary differential equations*, BIT, 40 (2000), pp. 241–266.
- [13] C. FULTON, *Titchmarsh–Weyl m -functions for second-order Sturm–Liouville problems with two singular endpoints*, Math. Nachr., 281 (2008), pp. 1418–1475.
- [14] C. FULTON, D. PEARSON, AND S. PRUESS, *Computing the spectral function for singular Sturm–Liouville problems*, J. Comput. Appl. Math., 176 (2005), pp. 131–162.
- [15] C. FULTON, D. PEARSON, AND S. PRUESS, *Efficient calculation of spectral density functions for specific classes of singular Sturm–Liouville problems*, J. Comput. Appl. Math., 212 (2008), pp. 150–178.
- [16] C. FULTON, D. PEARSON, AND S. PRUESS, *New characterizations of spectral density functions for singular Sturm–Liouville problems*, J. Comput. Appl. Math., 212 (2008), pp. 194–213.
- [17] C. FULTON, D. PEARSON, AND S. PRUESS, *Estimating spectral density functions for Sturm–Liouville problems with two singular endpoints*, arXiv:1303.2989, <http://arxiv.org/abs/1303.2989>, 2013.
- [18] C. FULTON AND S. PRUESS, *The computation of spectral density functions for singular Sturm–Liouville problems involving simple continuous spectra*, ACM Trans. Math. Software, 24 (1998), pp. 107–129.

- [19] C. FULTON, S. PRUESS, AND W. SHOAFF, *Parallel computation of Sturm–Liouville spectral density functions*, *Parallel Algorithms Appl.*, 4 (1994), pp. 41–51.
- [20] G.P. GHIROLDI AND L. GIBELLI, *A direct method for the Boltzmann equation based on a pseudo-spectral velocity space discretization*, *J. Comput. Phys.*, 258 (2014), pp. 568–584.
- [21] E. HAIRER, S.P. NORSETT, AND G. WANNER, *Solving Ordinary Differential Equations I: Non-stiff Problems*, 2nd ed., Springer, Berlin, 2000.
- [22] M. HAJMIRZAAHMAD AND A.M. KRALL, *Singular second-order operators: The maximal and minimal operators, and selfadjoint operators in between*, *SIAM Rev.*, 34 (1992), pp. 614–634.
- [23] R. HAMMERLING, O. KOCH, AND E.B. WEINMÜLLER, *Numerical solution of singular ODE eigenvalue problems in electronic structure computations*, *Computer Phys. Comm.*, 181 (2010), pp. 1557–1561.
- [24] R.D. HAZELTINE AND J.D. MEISS, *Plasma Confinement*, *Frontiers in Physics*, Addison–Wesley, Redwood City, CA, 1992.
- [25] R.D. HAZELTINE AND F.L. WAELBROECK, *The Framework of Plasma Physics*, Perseus, Reading, MA, 1998.
- [26] P. HELANDER, *Collisional Transport in Magnetized Plasmas*, Cambridge University Press, Cambridge, UK, 2002.
- [27] W.P. JOHNSON, *The curious history of Faá di Bruno’s formula*, *Amer. Math. Monthly*, 109 (2002), pp. 217–234.
- [28] M. LANDREMAN AND D.R. ERNST, *Local and global Fokker–Planck neoclassical calculations showing flow and bootstrap current modification in a pedestal*, *Plasma Phys. Control. Fusion*, 54 (2012), 115006.
- [29] M. LANDREMAN AND D.R. ERNST, *New velocity-space discretization for continuum kinetic calculations and Fokker–Planck collisions*, *J. Comput. Phys.*, 243 (2013), pp. 130–150.
- [30] R. LAVINE, *Absolute continuity of positive spectrum for Schrödinger operators with long-range potentials*, *J. Funct. Anal.*, 12 (1973), pp. 30–54.
- [31] V.L. MAKAROV, D.V. DRAGUNOV, AND Y.V. KLIMENKO, *The FD–method for solving Sturm–Liouville problems with special singular differential operator*, *Math. Comp.*, 82 (2013), pp. 953–973.
- [32] R. MILSON, *Liouville transformation and exactly solvable Schrödinger equations*, *Internat. J. Theoret. Phys.*, 37 (1998), pp. 1735–1752.
- [33] F.W.J. OLVER, *Asymptotics and Special Functions*, A. K. Peters, Wellesley, MA, 1997.
- [34] A. PATAKI AND L. GREENGARD, *Fast elliptic solvers in cylindrical coordinates and the Coulomb collision operator*, *J. Comput. Phys.*, 230 (2011), pp. 7840–7852.
- [35] D.B. PEARSON, *Singular continuous measures in scattering theory*, *Comm. Math. Phys.*, 60 (1978), pp. 13–36.
- [36] S. PRUESS AND C. FULTON, *Mathematical software for Sturm–Liouville problems*, *ACM Trans. Math. Software*, 19 (1993), pp. 360–376.
- [37] M. REED AND B. SIMON, *Functional Analysis*, Academic Press, San Diego, CA, 1980.
- [38] R.D. RICHTMYER, *Principles of Advanced Mathematical Physics*, Springer, New York, 1978.
- [39] M.N. ROSENBLUTH, W.M. MACDONALD, AND D.L. JUDD, *Fokker–Planck equation for an inverse-square force*, *Phys. Rev.*, 107 (1957), pp. 1–6.
- [40] W. RUDIN, *Real and Complex Analysis*, McGraw–Hill, New York, 1987.
- [41] B. SHIZGAL, *A Gaussian quadrature procedure for use in the solution of the Boltzmann equation and related problems*, *J. Comput. Phys.*, 41 (1981), pp. 309–328.
- [42] B. SIMON, *Orthogonal Polynomials on the Unit Circle, Part 1: Classical Theory*, AMS, Providence, RI, 2004.
- [43] I. STAKGOLD, *Green’s Functions and Boundary Value Problems*, Wiley, New York, 1998.
- [44] E.C. TITCHMARSH, *Eigenfunction Expansions Associated with Second-Order Differential Equations*, 2nd ed., Clarendon Press, Oxford, UK, 1962.
- [45] L.N. TREFETHEN, *Approximation Theory and Approximation Practice*, SIAM, Philadelphia, 2013.
- [46] J. WEIDMANN, *Zur spectraltheorie von Sturm–Liouville-operatoren*, *Math. Zeitschr.*, 98 (1967), pp. 268–302.
- [47] H. WEYL, *Über gewöhnliche Differentialgleichungen mit Singularitäten und die zugehörigen Entwicklungen willkürlicher Funktionen*, *Math. Ann.*, 68 (1910), pp. 220–269.
- [48] J. WILKENING, *High-order methods for computing spectral density functions for singular Sturm–Liouville problems*, in preparation.
- [49] J. WILKENING, A. CERFON, AND M. LANDREMAN, *Accurate spectral numerical schemes for kinetic equations with energy diffusion*, submitted; arXiv:1402.2971, <http://arxiv.org/abs/1402.2971>.