

Surprises in a Classic Boundary-Layer Problem*

William A. Clark[†]

Mario W. Gomes[‡]

Arnaldo Rodriguez-Gonzalez[§]

Leo C. Stein[¶]

Steven H. Strogatz[†]

Abstract. We revisit a textbook example of a singularly perturbed nonlinear boundary-value problem. Unexpectedly, it shows a wealth of phenomena that seem to have been overlooked previously, including a pitchfork bifurcation in the number of solutions as one varies the small parameter, and transcendentally small terms in the initial conditions that can be calculated by elementary means. Based on our own classroom experience, we believe this problem could provide an enjoyable workout for students in courses on perturbation methods, applied dynamical systems, or numerical analysis.

Key words. boundary layer, singular perturbation, dynamical systems, higher-order matching, matched asymptotics, shooting method

MSC codes. 34B16, 34E05, 34E15, 37M20, 65L11

DOI. 10.1137/21M1436087

I. Introduction. In many parts of mathematics, physics, engineering, and the life sciences, researchers have developed ingenious techniques for gaining insight into difficult problems by exploiting the presence of a small parameter in them. These techniques, known as perturbation methods, have shed light on all sorts of fascinating phenomena in fluid dynamics, mathematical biology, optics, chemical engineering, quantum mechanics, plasma physics, climate science, and many other disciplines. See [2, 11, 12, 16, 18, 19, 22] for just a few of the textbook introductions to perturbation methods and their applications.

This is the story of a textbook problem that has surprised us over and over again, and for which we have come to feel genuine affection. The problem is to solve the following nonlinear differential equation, subject to the given boundary conditions:

$$(1.1) \quad \epsilon y'' = yy' - y, \quad y(0) = 1, \quad y(1) = -1.$$

*Received by the editors July 26, 2021; accepted for publication (in revised form) March 15, 2022; published electronically February 9, 2023.

<https://doi.org/10.1137/21M1436087>

Funding: This work was supported by the NSF through grant DMS-1645643.

[†]Department of Mathematics, Cornell University, Ithaca, NY 14853 USA (wac76@cornell.edu; shs7@cornell.edu, stevenstrogatz.com).

[‡]Department of Mechanical Engineering, Rochester Institute of Technology, Rochester, NY 14623 USA (mwgame@rit.edu).

[§]Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853 USA (ajr295@cornell.edu).

[¶]Department of Physics and Astronomy, The University of Mississippi, University, MS 38677 USA (lcstein@olemiss.edu).

Here we use the notation $y' = dy/dx$, and we want to solve for $y(x)$ on the domain $0 \leq x \leq 1$ under the assumption that the parameter ϵ is small: $0 < \epsilon \ll 1$.

We first met problem (1.1) in the classic textbook by Mark Holmes [12]. Soon after that book appeared in 1995, one of us (Strogatz) decided to discuss (1.1) in an introductory course on asymptotics and perturbation methods. It quickly became clear the problem contains unexpected subtleties and riches.

In the years since then, whenever Strogatz has had a chance to teach that course, he has revisited problem (1.1) and has always learned something new about it, thanks to the questions and insights of the students, postdocs, and colleagues in attendance, most notably the coauthors on this paper. Together we think we may have finally gotten to the bottom of it.

But because this is meant to be an Education paper, we will also be pointing out some of the false turns we took along the way. Confusion is a natural part of doing mathematics. We did not land on the right way to think about (1.1) initially; it required lots of trial and error and a willingness to be open and vulnerable about what puzzled us at any given time. Experienced researchers know this is all part of the process, but we mention it for the sake of students who may be misled by the confident presentations and pristine appearance of the mathematics they see in most textbooks and journal articles. We want to be a little more honest here about how the sausage is actually made.

Our analyses have relied on three parts of applied mathematics: perturbation theory, nonlinear dynamics, and numerical analysis. As such, we think problem (1.1) could be useful to students or teachers in any of those subjects. We assume that the reader is comfortable with asymptotics and perturbation methods at the level of the books by Holmes [12] or Bender and Orszag [2], as well as nonlinear dynamics at the level of the book by Strogatz [24]. Not much exposure to numerical analysis is required; a basic knowledge of what it means to solve an ordinary differential equation numerically should be sufficient, say, at the level of someone who knows how to use *Mathematica* or *MATLAB* to solve an initial-value problem. For readers who want to immerse themselves in the details, we also provide supplementary notebooks: a *jupyter* notebook for the numerical methods and a *Mathematica* notebook for the analytical calculations [1].

Although we have grown enamored of problem (1.1) for its pedagogical value, it would be even more appealing if it also had some real-world applications. Alas, we have not found any so far, but it does have some close relatives of scientific interest. For example, if we change the sign of the right-hand side of the differential equation in problem (1.1), we get an equation known as the Lagerstrom–Cole equation [16, 19], which has been studied in aerodynamics and fluid dynamics as a model problem in connection with shock layers [19]. The differential equation in (1.1) also pops up in the study of delicate nonlinear phenomena known as canards [3, 7, 8, 9, 14, 17, 18]. As defined by Krupa and Szmolyan [17], “a canard solution is a solution of a singularly perturbed system which is contained in the intersection of an attracting slow manifold and a repelling slow manifold.” In applications, canards often arise in the analysis of nonlinear oscillations having both fast and slow time scales, as seen in chemistry, neurobiology, electronics, and many other fields [18].

2. The Surprises. So what are some of the surprises in problem (1.1)? The first is how many solutions it has. Holmes [12] argued that it has a unique solution, but it turns out that it actually has three. Figure 1 shows their graphs. We call these solutions B0, M, and B1, with the names chosen to indicate that they have a boundary




Fig. 1 Three numerical solutions of problem (1.1) for $\epsilon = 0.03$: two with boundary layers, and one with an interior layer. Under the symmetry $(x, y) \rightarrow (1-x, -y)$, the left and right boundary-layer solutions B0 and B1 trade places, and the interior-layer solution M is unchanged.

layer at $x = 0$, an interior layer in the middle, or a boundary layer at $x = 1$.

The graph of the B0 solution has a very negative initial slope $y'(0)$ at the left endpoint of the domain, while the other two solutions have $y'(0)$ *extremely* close to 1. In fact, those slopes differ from 1 by a “transcendently small” term, meaning a term that goes to zero faster than any positive power of ϵ as $\epsilon \rightarrow 0^+$. A common example is an exponentially small term of the form $\exp(-c/\epsilon)$, where $c > 0$ is a constant. Normally, such transcendently small terms are beyond the reach of basic perturbation theory (although sometimes they can be handled by more sophisticated techniques known as exponential asymptotics, superasymptotics, hyperasymptotics, or asymptotics beyond all orders [4, 5]). What especially surprised us about problem (1.1) is that the leading-order asymptotics of the initial slope $y'(0)$ could be found by elementary means, even when the difference between $y'(0)$ and 1 is transcendently small. As we show in section 6,

$$1 - y'(0) \sim \frac{9}{2\epsilon} \exp\left(-\frac{5}{8\epsilon}\right)$$

for the M solution, while for B1 it is given by

$$1 - y'(0) \sim \frac{24}{\epsilon} \exp\left(-\frac{3}{2\epsilon}\right).$$

These transcendently small terms can be calculated using nothing more than higher-order matching, phase plane analysis, and a constant of motion for the associated flow.

Not only were we surprised by this good fortune, we were also delighted and relieved by it, given that none of us knew anything about the more sophisticated techniques mentioned above!

Yet another surprise is the behavior of the solutions with respect to ϵ . In section 7 we'll see that as we increase ϵ from zero, the three solutions persist until a critical value $\epsilon_c \approx 0.216$, at which point they collide in a pitchfork bifurcation. For larger values of ϵ , problem (1.1) has a unique solution that resembles the M solution above.

3. Locating the Layers: Holmes's Analysis. Our initial thinking about problem (1.1) was greatly influenced by what Holmes [12] had to say about it. So we summarize his analysis here. We are not going to show the mathematical details yet; they will appear in later sections. For now we just want to emphasize a question that arises whenever one confronts a differential equation like (1.1) with a small parameter ϵ multiplying its highest derivative.

QUESTION 1: Does problem (1.1) have any layers (meaning regions of rapid variation in y or its derivatives), and if so, where are they located?

This issue comes up midway through Holmes's chapter on matched asymptotic expansions (Chapter 2 in [12]). Earlier in the chapter he has already introduced the basic ideas of inner and outer solutions, boundary layers, matching, and uniformly valid composite solutions. There is also a discussion of higher-order matching and examples with multiple boundary layers. In all these prior examples, the layers occur at the endpoints of the domain; in other words, they are genuine "boundary layers." In contrast, problem (1.1) is offered as the first instance of a problem having an "interior layer." Holmes writes:

Generally, when one first begins trying to solve a problem it is not known where the layer(s) is. If we began this problem as we did the previous two and assumed there is a boundary layer at either one of the endpoints, we would find that the expansions do not match. This is a lot of effort for no results, but fortunately there is a simpler way to come to the same conclusion.

He then offers a convexity argument to rule out boundary layers at either $x = 0$ or $x = 1$, but is careful to note that "these are only plausibility arguments and they do not prove anything. What they do is guide the analysis and hopefully reduce the work necessary to obtain the solution."

3.1. Holmes's Plausibility Argument. Holmes first considers a possible boundary layer at $x = 0$. He looks at the governing equation $\epsilon y'' = yy' - y$ along with its left boundary condition, $y(0) = 1$. Then he sketches the graph of a candidate solution that looks like our B0 solution in Figure 1, except that he also assumes that $y'' > 0$ in the boundary layer; in other words, $y(x)$ is assumed to be concave up for all small $x > 0$. On the other hand, although y'' (hypothetically) has one sign in the layer, y itself clearly changes sign from positive to negative as it drops from $y(0) = 1$ at the boundary to a value $y \approx -2$ where it matches the outer solution. With those sign considerations in mind, Holmes [12] continues:

If the solution behaves like the other example problems we have examined, then near $x = 0$ it would be expected that $y' < 0$ and $y'' > 0$. This means that $\epsilon y'' > 0$ but $y(y' - 1)$ is both positive and negative in the boundary layer. This is impossible It is possible to rule out a boundary layer at $x = 1$ in the same way.

Having ruled out (convex or concave) boundary layers at either end, Holmes [12]

next considers the possibility of an interior layer centered at a point $0 < x_0 < 1$. He works out three possible forms for the inner solution (depending on whether a certain constant of integration is positive, negative, or zero) and shows that only one of these can be matched to the outer solution. Then he uses a symmetry argument to conclude that the layer must be centered at the midpoint of the domain, $x_0 = 1/2$, and demonstrates that the inner solution must have odd symmetry about that point. Finally, he matches the inner and outer solutions, constructs a uniformly valid composite solution, and shows it agrees with a numerical solution obtained for $\epsilon = 0.01$. In this way, Holmes convincingly demonstrates the existence and properties of the solution we called M above, a solution of (1.1) with an interior layer in the middle of the domain.

But what about the possibility of nonconvex or nonconcave boundary-layer solutions? Recall that the plausibility argument only rules out solutions with y'' having strictly one sign in the layer. As Holmes's careful wording suggests, that loophole could potentially allow for sneakier solutions where y'' changes sign within a boundary layer. As we will see in the next section, that little finesse is precisely what allows the solutions B0 and B1 to exist.

4. Another Approach to Locating the Layers: Phase Plane Analysis. As we worked through Holmes's analytical approach to problem (1.1), we began to wonder if it might be helpful to supplement it with a more geometric style of reasoning known as phase plane analysis. After all, the equation appearing in (1.1), $\epsilon y'' = yy' - y$, is a nonlinear, second-order, autonomous differential equation, and phase plane analysis is a powerful tool for illuminating how the solutions to such equations behave. Plus, we have to admit, we have more experience with nonlinear dynamics than perturbation theory, so it felt like a more secure way to approach an unfamiliar problem.

4.1. Phase Portraits. To recast the problem into the language of dynamics, we replace the independent variable x with t , which we think of as time. Then the dependent variable y becomes a function of time t . The advantage of this approach is that it allows us to use our physical intuition about time and motion and the difference between fast versus slow. Abstract solutions to the differential equation turn into easily pictured trajectories of (imaginary) particles moving around in a two-dimensional space known as the phase plane.

To construct the phase plane, we convert the second-order equation $\epsilon y'' = yy' - y$ into a pair of first-order equations, which we then view as defining a vector field on the plane. We perform the first step by introducing a new dependent variable z , defined as $z = y'$, and then we rewrite $\epsilon y'' = yy' - y$ in terms of z to get $\epsilon z' = yz - y$. By solving for y' and z' and placing them on the left-hand side of a pair of first-order differential equations, we obtain the following vector field:

$$(4.1) \quad \begin{aligned} y' &= z, \\ z' &= \frac{1}{\epsilon}y(z - 1). \end{aligned}$$

Next we interpret (4.1) as a dynamical system. From this perspective, the vector (y', z') then tells us the instantaneous "velocity" that an imaginary particle at (y, z) would have at time t . As the imaginary particle moves around in the (y, z) phase plane, it traces out a trajectory $(y(t), z(t))$, which is the geometric counterpart of a solution $(y(x), y'(x))$ to the original problem (1.1).

This construction allows us to visualize how the solutions to (4.1) behave by imagining how particles move around in the phase plane. There is no need to be




Fig. 2 Phase portraits for (4.1) with $\epsilon = 0.1$. The left panel shows three qualitatively different kinds of trajectories, as well as fast and slow regions in the flow. An imaginary particle moves with $O(1)$ speed in the slow region (the thin gray strip around the invariant line $z = 1$). Outside the strip, the particle zips around much faster, with enormous vertical speeds of order $O(1/\epsilon)$. The right panel shows a quantitatively accurate phase portrait. Notice how tightly packed the trajectories become as they squeeze into and out of the slow region near $z = 1$.

quantitatively precise just yet; a qualitatively correct picture is enough at this stage. By looking at the signs of y' and z' in (4.1) and sketching a few vectors in various parts of the (y, z) plane, we are led to the picture shown in the left panel of Figure 2. This picture is called the phase portrait for the system. It shows that there are three qualitatively different types of trajectories for (4.1): parabolic-looking trajectories that flow from left to right above the horizontal line $z = 1$; a straight trajectory that flows from left to right along the invariant line $z = 1$; and periodic trajectories below $z = 1$ that form closed loops, and on which a particle would circulate round and round, always moving clockwise. The right panel of Figure 2 shows the same information quantitatively.

Incidentally, we can prove that the periodic-looking trajectories in Figure 2 are truly periodic and are not merely slowly winding spirals in disguise. There are two standard ways of proving this, so we will not dwell on the details; see sections 6.5 and 6.6 in [24] for an introduction. Briefly, one way is to note that (4.1) is a “conservative” system. To see this, rewrite it as $y'/z' = dy/dz = \epsilon z/[y(z-1)]$ and then separate variables and integrate to obtain $2\epsilon[z + \log(1-z)] - y^2 = \text{constant}$. This implicit equation can be shown to define closed curves for all $z < 1$. Another way is to observe that (4.1) is also a “reversible” system: the vector field (4.1) is unaltered by the change of variables $(x, y, z) \rightarrow (-x, -y, z)$, which corresponds to a time reversal combined with a mirror reflection across the z -axis in Figure 2. From this symmetry we can conclude that the trajectories lying below $z = 1$ are composed of two left/right mirror-image halves that together form a bilaterally symmetric loop.

4.2. An Important Symmetry. A stronger reversibility symmetry of (4.1) is worth noting: Both the differential equation $\epsilon y'' = yy' - y$ and its boundary conditions $y(0) = 1, y(1) = -1$ are left unchanged by the transformation

$$(x, y) \rightarrow (1 - x, -y).$$




Fig. 3 Phase plane plots of the three solutions of the boundary-value problem (1.1) for $\epsilon = 0.1$. The vertical coordinate z denotes y' . The red, green, and blue trajectories correspond to the B0, M, and B1 solutions having layers at $x = 0$, $x = 1/2$, and $x = 1$, respectively. The beginning ($x = 0$) of each trajectory is denoted with a circle, and the end ($x = 1$) is denoted with a square; the flow is clockwise on each trajectory. All trajectories start on the dashed vertical line $y = 1$ and end on the dashed vertical line $y = -1$, corresponding to the original boundary conditions. Note that the B0 and B1 trajectories form a symmetric pair. As such, both of them lie on the same periodic orbit in the phase plane. The slow part of each trajectory (the “outer solution”) occurs at the top, close to $z = 1$, while the much faster part (the “inner solution” in the layer) occurs everywhere else below that.

Hence the possible solutions of the original boundary-value problem (1.1) either come in symmetrical pairs $y(x)$ and $\tilde{y}(x)$, where $\tilde{y}(x) = -y(1-x)$, or the pair degenerates to a single solution with the symmetry $y(x) = -y(1-x)$. We saw a visual manifestation of this symmetry in Figure 1, where B0 and B1 form a symmetric pair and M is self-symmetric. Likewise, when those same solutions are plotted in the phase plane shown in Figure 3, the red and blue curves are paired under the symmetry and the green curve is self-symmetric.

4.3. Boundary Conditions. What about the boundary conditions $y(0) = 1$ and $y(1) = -1$ in (1.1)? How do they enter the phase plane picture? Well, the condition $y(0) = 1$ means that at time $t = 0$ our imaginary particle must start somewhere on the vertical line $y = 1$ in the phase plane (Figure 3). Its z coordinate on that line, however, is unspecified and remains to be determined; indeed, the key to solving (1.1) is to figure out the initial value of z that will enable the moving particle to satisfy the *other* boundary condition, $y(1) = -1$. In dynamical terms, this other boundary condition $y(1) = -1$ is a final condition, not an initial condition. It says that the particle must reach the vertical line $y = -1$ in the (y, z) plane after exactly one unit of travel time.

Thus, we see what a difficult challenge our imaginary particle is facing. It must find exactly the right place to start on the line $y = 1$, such that after it gets carried along by the flow determined by the vector field, somewhat like a tiny speck of leaf being carried downstream by a gentle brook, it manages to land somewhere on the line $y = -1$ precisely when the clock strikes time $t = 1$.

Figure 2 immediately implies that the trajectories on the line $z = 1$ or above it are disqualified as candidate solutions because there’s no way they can satisfy the boundary conditions: A particle starting on any one of them would move monotonically to

the right, and so would have no chance of making the leftward journey required to get from $y = 1$ to $y = -1$. That leaves the closed loops below $z = 1$ as our only hope. And indeed, we can imagine a particle starting somewhere on the line $y = 1$ and then flowing along an arc on one of the closed loops such that if everything is chosen just right (i.e., we pick the right loop to start on), the particle will reach $y = -1$ after exactly one unit of travel time. Figure 3 shows three arcs that do the trick.

4.4. Slow-Fast Structure. So far we have not used the assumption that ϵ is small, but now we will. For $\epsilon \ll 1$, we see from (4.1) that the vector field has large regions where the flow is very fast in the vertical direction, with vertical velocities z' of order $O(1/\epsilon)$ occurring at all points in the (y, z) plane where $z - 1 = O(1)$. This region of fast variation, as we will soon see, corresponds to the “inner region” in a perturbation treatment via boundary-layer theory. In the phase portrait, it consists of all points outside the thin gray strip shown schematically in Figure 2. The strip does not have well-defined edges, but its blurriness does not matter; the key thing is that its thickness is $O(\epsilon)$, however we define it. (To check its thickness, observe that if y is $O(1)$ and $z - 1 = O(\epsilon)$, then y' and z' are both $O(1)$ in (4.1), indicating that the flow is slow compared to the $O(1/\epsilon)$ speeds achieved everywhere outside the strip.) This strip in the phase plane where the motion is comparatively slow corresponds to the “outer region” in a boundary-layer treatment.

It was by contemplating the slow-fast structure of the flow that we originally came to suspect that there might be more than one solution to problem (1.1). As we reconsidered the solution discussed by Holmes, with its interior layer centered at $x = 1/2$, we pictured it as a particle moving with a slow-fast-slow trajectory in the phase plane, as shown in the middle panel of Figure 3. Our imaginary particle spends about half of its travel time dawdling through the initial slow region at the top of the green arc in Figure 3, then rockets down and around and up again through the fast region in almost no time at all, and then dawdles through the remaining slow region for the remaining half of its travel time. Why, we wondered, couldn't a particle spend nearly *all* its time in a slow region at the beginning? Or at the end? If the particle started sufficiently close to the invariant line $z = 1$, or ended up near there, it seemed like these sorts of solutions should also be possible.

This intuition turned out to be correct: such solutions do exist. The blue and red curves in Figure 3 show what they look like as trajectories.

4.5. More Back Story: Puzzling over the Initial Slope. It took us considerable trial and error to find these boundary-layer solutions numerically the first time we looked for them using the computer, more than a decade ago. They eluded us completely when ϵ was very small. Fortunately, for ϵ only moderately small it was not difficult to find them. For $\epsilon = 0.1$, for example, we found that $y'(0) \approx 0.9999$ yielded a trajectory that was slow at the beginning and fast at the end (the B1 solution), whereas $y'(0) \approx -10.6942$ gave a trajectory that was fast and then slow (the B0 solution).

The strikingly small difference between 0.9999 and 1 made us wonder what the formula for $y'(0)$ as a function of ϵ might be for the B1 solution. Likewise, given the size of -10.6942 , we were curious how negative $y'(0)$ might get for smaller values of ϵ as we continued tracking the B0 solution. Our numerics couldn't answer these questions at the time, since some of us were naive about computational methods, so we did not know how to solve the boundary-value problem reliably for $\epsilon \ll 1$.

Perhaps the initial slope could be found by asymptotic analysis? We felt sure it could but did not immediately see how to do it. We pose that as our next big question

and call it the *puzzle of the initial slope*.

QUESTION 2: For the three solutions of problem (1.1), how do their initial slopes $y'(0)$ depend on ϵ for $0 < \epsilon \ll 1$?

We will work toward answering that question in the next few sections. But before we leave the phase plane, we should notice that it tells us one more thing of interest. It reveals that the B0 and B1 solutions sneak through the loophole in Holmes's plausibility argument by having layers that are *nonconvex*. To see how that conclusion follows from the phase portrait, observe that on any of the trajectories shown in Figure 3, the value of z' (the vertical velocity) evidently *changes sign* as the particle goes down and then back up on its journey through the fast region. Since $z' = y''$, that change of sign means the concavity of $y(x)$ changes sign in the layer!

5. Perturbation Theory. In this section we solve problem (1.1) for $\epsilon \ll 1$, both inside and outside the boundary layers or interior layers. Then we match the inner and outer solutions and find composite solutions that give uniformly valid asymptotic approximations of $y(x)$ over the whole domain $0 \leq x \leq 1$. We perform the match to first order in ϵ (i.e., we go beyond the leading order of perturbation theory) because it turns out we need this higher-order information to solve the puzzle of the initial slope (Question 2). In that sense, the following analysis provides a motivational case study of why one would ever want to do higher-order matching. The details of this analysis are included in the supplementary *Mathematica* notebook [1].

5.1. Outer Solution. First we consider the outer region, where regular perturbation theory applies. In this region, we expand $y(x, \epsilon)$ in the regular perturbation series

$$(5.1) \quad y(x, \epsilon) \sim y_0(x) + \epsilon y_1(x) + O(\epsilon^2),$$

insert this into the differential equation $\epsilon y'' = yy' - y$ appearing in (1.1), and collect terms having like powers of ϵ . At leading order we find

$$(5.2) \quad y_0 y_0' - y_0 = 0.$$

This equation has two possible solutions: $y_0 = 0$ (which cannot satisfy the boundary conditions) or $y_0' = 1$, yielding $y_0(x) = x + a$ for some real constant a . In fact, a bit of study shows that the higher corrections all satisfy $y_n' = 0$ for $n > 0$, and thus $y_n = a_n$ for some constants a_n . But the leading-order solution $y_0(x)$ already fixes the constant by satisfying the boundary condition, so all the higher constants vanish: $a_n = 0$ for all $n > 0$. Therefore $y_0(x) = x + a$ is not merely the zeroth-order approximation to the outer solution; it is the outer solution at *all orders* of ϵ . We can also reach this conclusion by noting that $y_0(x) = x + a$ satisfies the original differential equation $\epsilon y'' = yy' - y$ exactly for all values of ϵ .

For an outer solution that includes $x = 0$ in its domain, we can determine the constant a by applying the boundary condition at that endpoint, and similarly for an outer solution that includes $x = 1$. These two potential outer solutions that satisfy either the left or right boundary condition are

$$(5.3) \quad y_0^L(x) = 1 + x, \quad y_0^R(x) = x - 2.$$

5.2. Inner Equation. Now we move on to the inner solutions. Suppose there is a layer at $x = x_0$. Holmes [12] shows that $x_0 = 1/2$ is the only possible location for

an interior layer, and we are going to show that boundary layers can occur at $x_0 = 0$ or $x_0 = 1$ as well. As before, we refer to these three cases as M (layer in the middle), B0 (layer on the boundary at $x = 0$), and B1 (layer on the boundary at $x = 1$).

Introduce a layer thickness δ , which is a function of ϵ to be determined. Let

$$(5.4) \quad X \equiv \frac{x - x_0}{\delta}$$

be a scaled independent variable that describes positions in the layer, and let

$$(5.5) \quad Y(X) \equiv y(x) = y(x_0 + \delta X)$$

be a new dependent variable that describes how y varies in the layer. Derivatives of the new variable are

$$(5.6) \quad \frac{dY}{dX} = \delta y', \quad \frac{d^2Y}{dX^2} = \delta^2 y''.$$

Now our original differential equation $\epsilon y'' = yy' - y$ becomes

$$(5.7) \quad \frac{\epsilon}{\delta^2} \frac{d^2Y}{dX^2} = Y \left(\frac{1}{\delta} \frac{dY}{dX} - 1 \right).$$

If δ is chosen correctly, then Y and all its derivatives should be $O(1)$ as $\epsilon \rightarrow 0$. Thus we find a distinguished limit when $\epsilon/\delta^2 = 1/\delta$, or simply $\delta = \epsilon$. Therefore the inner equation in the layer is given by

$$(5.8) \quad \frac{d^2Y}{dX^2} = Y \frac{dY}{dX} - \epsilon Y.$$

5.3. Leading-Order Inner Solution. To solve the inner equation (5.8) asymptotically, we expand Y as

$$(5.9) \quad Y(X, \epsilon) \sim Y_0(X) + \epsilon Y_1(X) + O(\epsilon^2).$$

Inserting this series into (5.8) yields, at leading order,

$$(5.10) \quad \frac{d^2Y_0}{dX^2} = Y_0 \frac{dY_0}{dX} = \frac{d}{dX} \left(\frac{1}{2} Y_0^2 \right).$$

This can be integrated to obtain

$$(5.11) \quad \frac{dY_0}{dX} = \frac{1}{2} Y_0^2 + A,$$

where A is an integration constant.

This result has a nice geometrical interpretation. If we recall that $y' = z$ in the phase plane, then (5.11) shows that, to leading order, the trajectories in the (y, z) plane follow *parabolic* arcs as they move through the inner region where the motion is fast. From the phase plane pictures shown earlier, we know that the only parabolic arcs of interest are those with a negative z -intercept, as these are the arcs that lie on the closed loops. Hence we see that A should be negative, say, $A = -\frac{1}{2}b^2$.

Then separating the variables in (5.11) and integrating gives

$$(5.12) \quad \frac{dY_0}{Y_0^2 - b^2} = \frac{1}{2} dX,$$

$$(5.13) \quad Y_0(X) = b \tanh \left(c - \frac{b}{2} X \right),$$

with an additional integration constant c . The two constants b, c are determined by matching to the solution in the outer region.

5.4. Zeroth-Order Matching for the Symmetric Solution M . Let's see how the matching goes for the interior layer at $x_0 = 1/2$. We know we have a symmetric solution that satisfies $y^M(1/2) = 0 = Y_0^M(0)$ (we could also learn this from the matching alone). This symmetry condition tells us that Y_0^M is an odd function of X and hence $c = 0$. To find the value of b , we need to look at the large- X asymptotic behavior of the inner solution and match it to the asymptotic behavior of the outer solution as x approaches $x_0 = 1/2$ from either side. Thus we need to take the limit $\lim_{X \rightarrow \pm\infty} Y^M(X)$ and match it to $y_0^L(1/2) = 3/2$ and $y_0^R(1/2) = -3/2$. If we recall that $\tanh z$ is odd and $\lim_{z \rightarrow +\infty} \tanh z = 1$, we see that both limits agree if and only if $b = \pm 3/2$. For either choice of b , our inner solution at zeroth order becomes

$$(5.14) \quad Y_0^M = -\frac{3}{2} \tanh\left(\frac{3}{4}X\right).$$

Finally, we can construct a composite solution y_c by the usual recipe: $y_c = y_{\text{outer}} + Y_{\text{inner}} - y_{\text{match}}$. Carrying out those steps for the zeroth-order approximation to M yields

$$(5.15) \quad y_{c,0}^M = x - \frac{1}{2} - \frac{3}{2} \tanh\left[\frac{3}{4\epsilon}\left(x - \frac{1}{2}\right)\right] + O(\epsilon).$$

5.5. Zeroth-Order Matching for the Asymmetric Solutions B_0 and B_1 . We can proceed similarly for the B_0 and B_1 cases. In fact, we only need to do the work for one of them, since we can get the other one from the symmetry transformation $(x, y) \rightarrow (1 - x, -y)$. Let's focus on the B_0 case, which has a layer at $x_0 = 0$. Its zeroth-order inner solution (5.13) needs to be matched to $y_0^R(0) = -2$, which tells us that $b = -2$. We also need to satisfy the boundary condition $y(0) = 1$, which is now inside the layer, so $Y_0^{B_0}(0) = 1$ determines the value of c as $c = -\tanh^{-1} \frac{1}{2}$. Hence

$$(5.16) \quad Y_0^{B_0} = -2 \tanh\left(X - \tanh^{-1} \frac{1}{2}\right).$$

Similarly, for B_1 we get

$$(5.17) \quad Y_0^{B_1} = -2 \tanh\left(X + \tanh^{-1} \frac{1}{2}\right).$$

Now, constructing the leading-order composite solutions, we get

$$(5.18) \quad y_{c,0}^{B_0} = x - 2 \tanh\left(\frac{x}{\epsilon} - \tanh^{-1} \frac{1}{2}\right) + O(\epsilon),$$

$$(5.19) \quad y_{c,0}^{B_1} = x - 1 - 2 \tanh\left(\frac{x - 1}{\epsilon} + \tanh^{-1} \frac{1}{2}\right) + O(\epsilon).$$

5.6. Remarks about the Zeroth-Order Composite Solutions. The leading-order composite solutions for M , B_0 , and B_1 look extremely similar to the numerical solutions plotted in Figure 1. They are also uniformly valid, as can be checked by looking at the difference between the numerical and analytical solutions, as plotted in Figure 4 for B_0 . Moreover, the layers for all of the zeroth-order solutions are nonconvex, as we expected from our earlier phase plane analysis. For example, the B_0 solution (5.18) has an inflection point at $x = \epsilon \tanh^{-1} \frac{1}{2}$.

Yet informative as these leading-order solutions are, they are not accurate enough to allow us to calculate the initial slope $y'(0)$ correctly. So, let's proceed to the next order.




Fig. 4 Scaled error in the composite zeroth-order asymptotic solution (5.18) for the case with a boundary layer at $x = 0$. The error is defined as the difference between the asymptotic solution (5.18) and a very careful numerical solution (taken as a surrogate for the unknown exact solution). The scaled error is defined here as the error divided by ϵ ; this scaling is appropriate because we expect to incur errors of size $O(\epsilon)$ in a leading-order solution. We examine the error within the boundary layer, $x = O(\epsilon)$, i.e., the layer has a width linear in ϵ (the error outside of the layer is transcendentally small). Notice that the composite solution is uniformly valid and the error is proportional to ϵ , as is expected from a zeroth-order solution. The error behavior is similar for the other two asymptotic solutions with layers at $x = 1/2$ and $x = 1$.

5.7. First-Order Matching for the Symmetric Solution M . As we noted earlier, the outer solutions are $y_0(x) = x + a$ to all orders in ϵ ; only the inner solution and the matching change as we proceed to higher orders. To study the first-order correction ϵY_1 in the inner solution, we insert the series (5.9) into our inner equation (5.8) and collect the first-order terms. We find that $Y_1(X)$ satisfies

$$(5.20) \quad \frac{d^2 Y_1}{dX^2} - \frac{dY_1}{dX} Y_0 - \frac{dY_0}{dX} Y_1 + Y_0 = 0.$$

Compared to the asymmetric B0 and B1 solutions, the symmetric solution M yields the simplest expressions for Y_1 , so we focus on that calculation now and relegate the others to Appendix A.

Using the zeroth-order solution for Y_0^M from (5.14), we find that Y_1^M satisfies the following second-order inhomogeneous linear differential equation with variable coefficients:

$$(5.21) \quad \frac{d^2 Y_1^M}{dX^2} + \frac{3}{2} \tanh\left(\frac{3X}{4}\right) \frac{dY_1^M}{dX} - \frac{9}{8} \operatorname{sech}^2\left(\frac{3X}{4}\right) Y_1^M = \frac{3}{2} \tanh\left(\frac{3X}{4}\right).$$

We must solve this beast subject to the side condition that $Y_1^M(0) = 0$ (a condition that follows from symmetry, as Y^M is an odd function). Impressively, **Mathematica** obliges and produces a long expression that we give in (A.1) of Appendix A. Of

the two free constants of integration appearing in (A.1), we fix the constant c_2 by imposing the side condition $Y_1^M(0) = 0$, giving $c_2 = \pi^2/18$. To determine c_1 , we need to perform a first-order match to the outer solution y_0^R on the right (there's no need to worry about additionally matching to y_0^L , the outer solution on the left; matching on the right automatically takes care of matching on the left, by the odd symmetry noted above).

To perform the matching on the right, we need to know the asymptotic behavior of Y_1^M as $X \rightarrow +\infty$. The relevant asymptotics are

$$(5.22) \quad \sinh z \sim \cosh z \sim \frac{1}{2}e^z \quad (z \rightarrow +\infty),$$

$$(5.23) \quad \log(1+z) \sim z \quad (z \rightarrow 0),$$

$$(5.24) \quad \text{Li}_s(z) \sim z \quad (z \rightarrow 0),$$

where $\text{Li}_s(z)$ is a special function called a ‘‘polylogarithm’’ of order s [21]. The polylogarithms can be defined by their power series or recursively from an integral:

$$(5.25) \quad \text{Li}_s(z) = \sum_{k=1}^{\infty} \frac{z^k}{k^s} = \int_0^z \text{Li}_{s-1}(t) \frac{dt}{t}.$$

For $s < 2$, they are elementary functions; for example, $\text{Li}_1(z) = -\log(1-z)$.

With further help from *Mathematica*, we eventually find

$$(5.26) \quad Y_1^M \sim X + \frac{2}{3}(c_1 - 1 - \log 4) \quad (X \rightarrow \infty).$$

Miraculously—and yet not miraculously at all, if one believes in perturbation theory—this function has exactly the right large- X behavior, $Y_1^M \sim X$, needed to match onto the *outer* solution, if the constant is correct! Recall that at zeroth order, $\lim_{X \rightarrow \infty} Y_0^M(X) = -3/2$ already matches the value of $\lim_{x \rightarrow 1/2^+} y_0^R(x)$. Therefore we want the additive constant above to vanish (or else we would make an $O(\epsilon)$ error in the matching). Hence

$$(5.27) \quad c_1 = 1 + \log 4.$$

Finally, by again invoking the recipe $y_c = y_{\text{outer}} + Y_{\text{inner}} - y_{\text{match}}$ for forming a composite solution, we obtain the composite solution up to terms of order ϵ^2 . Remarkably, the y_{match} that we need to subtract here is simply the exact outer solution y_0^R . So the whole composite solution boils down to the inner solution $Y_0^M + \epsilon Y_1^M$. Thus, the first-order composite solution for M is

$$(5.28) \quad \begin{aligned} y_{c,1}^M = & -\frac{3}{2} \tanh\left(\frac{3X}{4}\right) \\ & + \frac{\epsilon}{72} \text{sech}^2\left(\frac{3X}{4}\right) \left\{ 4\pi^2 + 48\text{Li}_2\left(-e^{-3X/2}\right) + 9X(8+3X) \right. \\ & \left. + 48 \sinh\left(\frac{3X}{2}\right) \log\left(2 \cosh\left(\frac{3X}{4}\right)\right) \right\} + O(\epsilon^2), \end{aligned}$$

where $X = (x - \frac{1}{2})/\epsilon$.

We proceed similarly for the B0 and B1 solutions, collecting the results in Appendix A. Figure 5 confirms that the error between the first-order asymptotic solution




Fig. 5 Scaled error, now defined as error divided by ϵ^2 , in the composite first-order asymptotic solution for the case with a boundary layer at $x = 0$, given in Appendix A. Notice that the composite solution is uniformly valid and the error is now proportional to ϵ^2 , as is expected for an asymptotic solution matched to first order.

for B0 and a numerical solution (presumed to be close to exact) truly does shrink in proportion to ϵ^2 , as it should for a first-order match. This sort of test provides a reassuring check when doing complicated numerics and asymptotics.

6. Solving for the Initial Slope. Now that we have constructed asymptotic approximations to the three solutions M, B0, and B1 of our original problem (1.1), we can use those approximations to estimate the initial slope $y'(0)$ in each case. Knowing this initial slope is theoretically interesting since (as we'll see) it depends on ϵ in an intriguing way. But it's also practically useful information: having a good approximation of the initial slope helps us solve the boundary-value problem (1.1) numerically.

One computational approach to solving a boundary-value problem is called the “shooting method.” A more thorough discussion can be found in several textbooks on numerical methods, for example, [23]. We also provide a supplementary `jupyter` notebook which implements the numerics described in this section [1].

Figure 6 shows an example of how shooting applies to our problem. We start a trial solution at $y(0) = 1$ and launch it with some initial slope $y'(0)$, somewhat like shooting an artillery shell at an intended target. In our case, the target is the point at the other boundary condition: $x = 1, y = -1$. Incorrect choices of the initial slopes at $x = 0$ will produce solutions that fail to hit the boundary condition at $x = 1$. The three curves shown on the left in Figure 6 indicate what happens if we aim too high or too low or just right.

If we compute where we hit for *many* possible $y'(0)$'s and plot the resulting $y(1)$'s versus $y'(0)$, we get the graph of the target function shown in the right panel of Figure 6. To find a numerical solution of our problem, then, we just have to figure out where the graph crosses the dashed horizontal line $y = -1$. This is a standard numerical task; it amounts to a root-finding problem in a small neighborhood of the

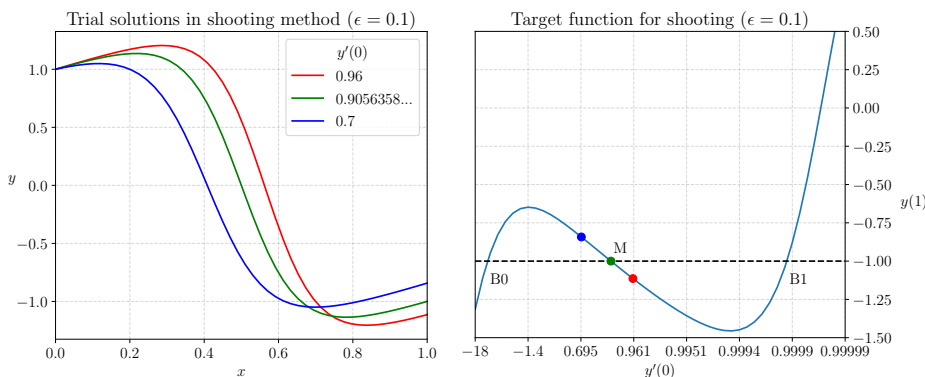


Fig. 6 Solving the boundary-value problem via shooting. Different initial slopes lead to different final values of y . The red, green, and blue trial solutions in the left panel have certain initial slopes $y'(0)$ and final $y(1)$ which correspond to the three filled circles plotted in the right panel. If we were to vary the initial slope continuously over a broader range, we would build up the curve in the right panel. Note that the horizontal axis in the right panel is uniform in $\log(1 - y')$, to expand the exponentially bunched region near $y'(0) = 1$. The initial slopes for the B0, M, and B1 solutions are marked with letters in the right panel; they occur at the three values of $y'(0)$ where the graph of the target function crosses the dashed horizontal line $y(1) = -1$.

solution. Notice that for the value $\epsilon = 0.1$ used to make Figure 6, the graph of the target function crosses the dashed line in three places. Those are the desired initial slopes of our three solutions.

To find analytical estimates of these slopes, we will see next why we needed to go to the trouble of doing higher-order matching.

6.1. Initial Slope for B0. Let's start by evaluating the value and slope of the zeroth-order composite solution for B0:

$$(6.1) \quad y_{c,0}^{B0}(0) = 1, \quad y_{c,0}^{B0'}(0) = -\frac{3}{2\epsilon} + 1.$$

The leading behavior here is $y'_{B0} \sim -3/(2\epsilon)$, and this is indeed correct, as we will soon see. However, the slope in (6.1) has an error that is $O(1)$. This can be seen graphically in Figure 4: Within the layer of width $O(\epsilon)$, there is an $O(\epsilon)$ error in the value of y , leading to an $O(1)$ error in the slope. Or, reiterating the point in another way, the derivative of a function's asymptotic series (at some order) is not necessarily the asymptotic series of the derivative of the function (to the same order). Since we have the composite solution to higher order, we can easily check what the correction is by using the results for B0 in Appendix A. Evaluation and differentiation of the first-order composite solution at $x = 0$ yield

$$(6.2) \quad y_{c,1}^{B0}(0) = 1, \quad y_{c,1}^{B0'}(0) = -\frac{3}{2\epsilon} + 1 + \log 16 + O(\epsilon).$$

This is now the full result for $y'_{B0}(0)$ up to errors of order $O(\epsilon)$. As shown by the top curve of Figure 7, this asymptotic result agrees nicely with the value of $y'(0)$ produced by the shooting method.

6.2. Initial Slope: The Tempting but Wrong Way. For the B0 case, we saw that the $O(1)$ term in the initial slope changed from 1 to $1 + \log 16$ when we went to first order. This change in the $O(1)$ term is a rather small effect as $\epsilon \rightarrow 0$, since the $-3/(2\epsilon)$ term dominates in this limit anyway. But to our surprise (probably because of our inexperience in these matters), we soon discovered that the B1 and M cases are much more subtle: differentiating the first-order composite solution does not give the correct initial slope, not even at leading order in ϵ .

It's an instructive trap to fall into, so let's take the plunge. For simplicity, let's work with M. If we calculate the initial value and initial slope of the zeroth-order composite solution $y_{c,0}^M$ given in (5.15), we find

$$(6.3) \quad y_{c,0}^M(0) = -\frac{1}{2} + \frac{3}{2} \tanh \frac{3}{8\epsilon} \sim 1 - 3e^{-\frac{3}{4\epsilon}} + \text{TST},$$

$$(6.4) \quad y_{c,0}^{M'}(0) = 1 - \frac{9}{8\epsilon} \operatorname{sech}^2 \frac{3}{8\epsilon} \sim 1 - \frac{9}{2\epsilon} e^{-\frac{3}{4\epsilon}} + \text{TST},$$

where TST stands for “transcendentally small terms,” i.e., terms that are smaller than any power of ϵ times the smallest reported term.

Should we trust these results? Let's check by going to the next order of perturbation theory. Our first-order composite solution yields

$$(6.5) \quad y_{c,1}^M(0) \sim 1 - e^{-\frac{3}{4\epsilon}} \left[\frac{3}{8\epsilon} + 4 + O(\epsilon) \right] + \text{TST},$$

$$(6.6) \quad y_{c,1}^{M'}(0) \sim 1 - e^{-\frac{3}{4\epsilon}} \left[\frac{9}{16\epsilon^2} + \frac{9}{2\epsilon} + \frac{\pi^2}{3} + O(\epsilon^2) \right] + \text{TST}.$$

This doesn't look good at all! Going to the next order in y has resulted in a change at a *lower* order to the values of $y(0)$ and $y'(0)$. Apparently we can't trust this. So for this problem at least, naively differentiating the composite solution does not give us the correct initial slope.

6.3. Initial Slope: The Right Way. Instead of the approach above, let's turn to a more global analysis, using our knowledge of the structure of solutions in the phase plane. The key insight is that we can use a conserved quantity (also known as a constant of motion, or a first integral) to transfer trustworthy information from inside a layer to a distant point outside the layer where we want to calculate an initial slope. For example, we can transfer information from the M layer at $x_0 = 1/2$ all the way over to $x = 0$; this trick is how we are going to extract the leading (but still minuscule!) transcendentally small term in M's initial slope. The advantage of using a conserved quantity is that it is exact; it allows us to shuttle information around in the phase plane with perfect fidelity.

6.3.1. Conserved Quantity for the Trajectories. To set the stage to perform the desired transfers of information, we need to gather a few facts about the conserved quantity.

We have already mentioned that our vector field (4.1) is conservative. Recall that at each point in the (y, z) phase plane,

$$(6.7) \quad \frac{dz}{dy} = \frac{z'(t)}{y'(t)} = \frac{y(z-1)}{\epsilon z},$$

which can be separated and integrated to yield

$$(6.8) \quad 2\epsilon[z + \log(1-z)] = y^2 - C^2,$$

where C is a constant that labels the trajectories. Since C is constant, the quantity

$$2\epsilon[z + \log(1 - z)] - y^2$$

remains unchanged as $y(t)$ and $z(t)$ flow along a trajectory and hence is a “conserved quantity.”

For a given value of C , we can generate two explicit formulas for the trajectories as curves in the (y, z) plane. Either we can write $y(z)$ in the right or left half plane by solving for y and using the relevant branch of the square root, or we can solve for $z(y)$ if we allow ourselves to use the implicitly defined Lambert W function, which satisfies

$$(6.9) \quad W(z) \exp(W(z)) = z.$$

This equation has multiple solutions, giving the multiple branches $W_n(z)$. For more on the Lambert W function, see [6]. From (6.8) the explicit solution for z is

$$(6.10) \quad z = 1 + W_n \left(-\exp \left(-1 + \frac{y^2 - C^2}{2\epsilon} \right) \right).$$

For $0 \leq z < 1$, we want the branch W_0 . For $z \leq 0$, we want the branch W_{-1} .

The existence of a conserved quantity for (4.1) suggests that the dynamical system might actually be a Hamiltonian system in appropriate coordinates. Indeed, (4.1) is Hamiltonian for $z < 1$: Canonical variables are $Q = y$ and $P = \log(1 - z)$, in terms of which the vector field corresponding to (4.1) becomes $Q' = 1 - e^P$, $P' = Q/\epsilon$, with a corresponding Hamiltonian $H(P, Q) = P - e^P - Q^2/(2\epsilon)$.

6.3.2. Using the Conserved Quantity to Transfer Information about Slopes.

Having set the stage, we're now ready to explain how to transfer slope information reliably with the help of the conserved quantity $2\epsilon[z + \log(1 - z)] - y^2$. Here's the idea: Recall that, by definition, $z = y'$, so z represents a slope y' on a graph of y versus x . Suppose we have one point (y_1, z_1) on a solution where we trust the slope z_1 . Then, by using the constancy of C , we can use this information to get the slope z_0 at some other y_0 on the same solution. For the M and B1 solutions, we will take (y_1, z_1) to be a convenient point inside the layer where the inner solution is trustworthy and then we can compute $z_1 = O(\epsilon^{-1})$ accurately. Then we will transfer this information over to the initial condition at $x = 0$, where $y_0 = 1$, and seek z_0 .

Taking two copies of (6.8) on the same curve of constant C and eliminating C , we find that our two points are related by

$$(6.11) \quad 2\epsilon[z_0 + \log(1 - z_0)] - y_0^2 = 2\epsilon[z_1 + \log(1 - z_1)] - y_1^2.$$

Let $y_0 = 1$. Solve for the desired initial slope z_0 by using the appropriate branch of the Lambert W function:

$$(6.12) \quad 1 - z_0 = -W_0 \left(-[1 - z_1] \exp \left(\frac{1 - y_1^2}{2\epsilon} - (1 - z_1) \right) \right).$$

Equation (6.12) is the key formula for transferring information from one point to another.

Now if we can find trustworthy values of y_1 and z_1 to plug in, we'll be in business. Let's illustrate the idea with the B1 solution. Remember that B1 is related to the B0

solution by the symmetry transformation $(x, y, z) \rightarrow (1 - x, -y, z)$. That means that the *final* slope in B1's layer at $x = 1$ is the same as the *initial* slope in B0's layer at $x = 0$. But we have already calculated B0's initial slope reliably! It's given in (6.2). (It's reliable because it was calculated inside a layer, namely, the layer at $x = 0$ for the B0 solution.)

Hence the trustworthy point (y_1, z_1) to pick for the B1 solution is

$$(6.13) \quad y_{B1}(1) = -1, \quad z_{B1}(1) = -\frac{3}{2\epsilon} + 1 + \log 16 + O(\epsilon).$$

When we plug this point into (6.12), we soon discover something fascinating about the initial slope $z_{B1}(0)$: it deviates from 1 by a transcendently small quantity. To see this, however, takes a few more steps. Upon performing the substitution we first obtain

$$(6.14) \quad 1 - z_{B1}(0) = -W_0 \left(\left[-\frac{3}{2\epsilon} + \log 16 + O(\epsilon) \right] \exp \left(-\frac{3}{2\epsilon} + \log 16 + O(\epsilon) \right) \right),$$

which may look opaque to anyone unfamiliar with the Lambert W function. Fortunately this expression can be simplified by using the fact that $W_0(z)$ is analytic at $z = 0$ and its power series is convergent within a radius $|z| < 1/e$ (see [6] for a thorough treatment). The first few terms are

$$(6.15) \quad W_0(z) = z - z^2 + O(z^3),$$

as can be verified by substituting back into the defining equation (6.9). We only need to keep the first term since the second term is already transcendently small in ϵ relative to the first term, in light of the form of the argument of W_0 in (6.14).

Thus, we finally arrive at the correct asymptotic behavior for B1's initial slope,

$$(6.16) \quad 1 - y'_{B1}(0) = \left[\frac{3}{2\epsilon} - \log 16 + O(\epsilon) \right] \exp \left(-\frac{3}{2\epsilon} + \log 16 + O(\epsilon) \right) + \text{TST},$$

which can be further simplified to

$$(6.17) \quad 1 - y'_{B1}(0) = \left[\frac{24}{\epsilon} + O(1) \right] \exp \left(-\frac{3}{2\epsilon} \right) + \text{TST}.$$

To obtain this last equality, we simplified (6.16) by replacing $\exp(\log 16)$ with 16 and $\exp(O(\epsilon))$ with $1 + O(\epsilon)$. That relative error at $O(\epsilon)$ in (6.16) times the leading ϵ^{-1} term results in our ignorance of the subdominant $O(1)$ term in the prefactor multiplying the controlling exponential in (6.17). Nevertheless, we still have enough information to nail down the leading-order ϵ -dependence of the initial slope, given by the term $[24/\epsilon] \exp(-3/2\epsilon)$.

In retrospect, these error considerations clarify *why* we needed go to the bother of approximating the inner solution with higher-order perturbation theory: It's because we needed the error in the argument of the exponential in (6.16) to be $O(\epsilon)$, which necessitated an error of $O(\epsilon^2)$ in the inner solution.

For the M solution, we redo the calculation above, except now we use the slope in the layer at $x = 1/2$ as our trustworthy value of z_1 . We get that slope with sufficient precision by simply evaluating the derivative of the first-order composite solution (5.28) at $x = 1/2$. We also recall that y vanishes at $x = 1/2$, by the odd symmetry of M about its midpoint. Hence the trustworthy values of y and its slope are

$$(6.18) \quad y_M(1/2) = 0, \quad y'_M(1/2) = -\frac{9}{8\epsilon} + 1 + \log 4 + O(\epsilon).$$




Fig. 7 Comparison of the initial slopes $y'(0)$ determined from the asymptotic solutions (dashed curves) and numerical solutions found via the shooting method (solid curves). The vertical axis plots $1 - y'(0)$ to highlight the tiny (transcendentally small) deviation of the initial slope from 1 for two of the solutions. From top to bottom, the curves give the behavior of the solutions labeled B0 (boundary layer at $x = 0$), M (interior layer at $x = 1/2$), and B1 (boundary layer at $x = 1$). Notice the vast difference in scale for the three initial slopes.

Plugging this phase space point into (6.12), we find

$$(6.19) \quad 1 - y'_M(0) = -W_0 \left(\left[-\frac{9}{8\epsilon} + \log 4 + O(\epsilon) \right] \exp \left(-\frac{5}{8\epsilon} + \log 4 + O(\epsilon) \right) \right)$$

$$(6.20) \quad = \left[\frac{9}{8\epsilon} - \log 4 + O(\epsilon) \right] \exp \left(-\frac{5}{8\epsilon} + \log 4 + O(\epsilon) \right) + \text{TST}$$

$$(6.21) \quad = \left[\frac{9}{2\epsilon} + O(1) \right] \exp \left(-\frac{5}{8\epsilon} \right) + \text{TST}.$$

Here, as in (6.16), we have enough precision to nail the leading-order term after expanding $\exp(O(\epsilon))$, but not enough to determine the subleading $O(1)$ correction.

Figure 7 compares these analytical initial slopes against numerical results obtained from the shooting method. In all three cases, the agreement between asymptotics and numerics is excellent. And for both M and B1, the agreement extends over more than twenty orders of magnitude. Wow!

7. Pitchfork Bifurcation. As we've seen, the nonlinear boundary-value problem (1.1) has three solutions when ϵ is sufficiently small. The final surprise in this problem comes when we examine what happens to the three corresponding initial slopes, $y'(0)$, as we increase ϵ away from 0. We can mull over the three equations (6.2), (6.16), and (6.21), or examine Figure 7 that summarizes all three. For sufficiently small ϵ , we have the ordering $y'_{B0}(0) < y'_M(0) < y'_{B1}(0)$. As ϵ increases, $y'_{B0}(0)$ increases (becomes less negative) and moves toward zero, while both $y'_M(0)$ and $y'_{B1}(0)$ decrease away from 1 and also move toward zero.

This behavior suggests that at sufficiently large ϵ , there is a possibility that two




Fig. 8 A pitchfork bifurcation occurs in the space of solutions to the boundary-value problem (1.1). Each point corresponds to integrating the initial-value problem with $y(0) = 1$, $y'(0)$ determined by the vertical coordinate on the plot, at a certain value of ϵ determined by the horizontal coordinate on the plot. The color denotes the value of $1 + y(1)$, with reds denoting positive values and blues denoting negative values. The curves where $1 + y(1) = 0$ lie between the reds and blues; these curves define the set of all $(\epsilon, y'(0))$ that yield solutions of the boundary-value problem. For each $\epsilon < \epsilon_c$, there are three initial values satisfying $1 + y(1) = 0$, corresponding to the B0, M, and B1 solutions, which coalesce at $\epsilon_c = 0.2159869288903\dots$ and $y'(0) = 0$. The zoom in the lower panel expands the neighborhood of the bifurcation point, showing that it has the form of a cubic, as is expected from the normal form of a pitchfork bifurcation.

or even three solutions might approach each other and merge. It also suggests that the initial slopes of the merging solutions might lie somewhere close to $y'(0) = 0$. We first discovered experimentally that this actually happens; a three-way merger occurs through a pitchfork bifurcation precisely when $y'(0) = 0$. After the fact, we were able to establish the analytical conditions that describe the bifurcation and identify the critical value ϵ_c where it occurs.

7.1. Visualizing and Explaining the Pitchfork Bifurcation. We visualize the merger of solutions by plotting contours of $1 + y(1)$ in the $(\epsilon, y'(0))$ plane in Figure 8. Satisfying the boundary-value problem means finding the contours where $y(1) + 1 = 0$, which occurs at the boundary between reds and blues in the figure. Below the critical value $\epsilon < \epsilon_c$, there are three solutions; these bifurcate at $\epsilon = \epsilon_c$ and $y'(0) = 0$. For large values of ϵ , there is only one solution. In the broad view shown in the top panel,

there is a rather obscure shape for the solution set. In the bottom panel, we zoom in to the sixth decimal place in ϵ and see a classic pitchfork shape, which we will explain below.

The first condition for the bifurcation comes from understanding the relationship between the initial and final slopes for the B0 and B1 solutions, which are paired under the symmetry $(x, y, z) \rightarrow (1 - x, -y, z)$. Let's return to (6.11), the relationship between two points (y_0, z_0) and (y_1, z_1) on the same trajectory. Take these two points to be the endpoints of the trajectory, where $y_0 = 1$ and $y_1 = -1$. Then the y^2 terms cancel, leaving a simpler condition relating the two endpoint slopes:

$$(7.1) \quad z_0 + \log(1 - z_0) = z_1 + \log(1 - z_1).$$

For convenience, let us define a function

$$(7.2) \quad f(z) = z + \log(1 - z),$$

whose domain is $z < 1$. Equation (7.1) says that $f(z_0) = f(z_1)$. Now there are two possibilities. For the self-symmetric solution M, the slopes at the endpoints agree automatically, so $z_0 = z_1$ and (7.1) is vacuously satisfied. This case tells us nothing new. However, for the asymmetric solutions B0 and B1, there are two distinct slopes, $z_0 \neq z_1$, and we can show that they must have opposite signs.

This result follows from the shape of the graph of f . Note first that $f'(z) = 1 - 1/(1 - z)$, so f has only one local extremum, at $z = 0$, where $f(0) = 0$. The second derivative $f''(z) = -(1 - z)^{-2}$ is strictly negative everywhere, so $z = 0$ is in fact a global maximum. For $z < 0$, $f(z)$ is monotonically increasing, while for $0 < z < 1$, $f(z)$ is monotonically decreasing. Therefore, in the case where we have two distinct slopes $f(z_0) = f(z_1)$ but $z_0 \neq z_1$, we can deduce that they must lie in the intervals $z_0 < 0$ and $0 < z_1 < 1$, or vice versa, and hence have opposite signs, as claimed.

Graphically, every horizontal slice through the graph of $f(z)$ below its global maximum cuts it in two places: one value at $z_0 < 0$ and the other value at $0 < z_1 < 1$. As we increase our horizontal slice toward the maximum, these two roots both approach and ultimately coalesce at $z = 0$.

Finally, let's apply this knowledge to the slopes z_{B0} and z_{B1} , when both solutions exist. The B0 solution has endpoint slopes $y'(0) = z_{B0}$ and $y'(1) = z_{B1}$, and the B1 solution exchanges these two. By the previous analysis, $f(z_{B0}) = f(z_{B1})$ with $z_{B0} < 0$, while $0 < z_{B1} < 1$. We know that as ϵ increases, these two slopes approach one another. The only place they can merge is at $z_{B0} = z_{B1} = 0$, which happens at some $\epsilon = \epsilon_c$.

What about the intermediate M solution? How does its slope compare to the other two? We know that the three slopes are the three roots of the target function plotted in the right panel of Figure 6. Is it possible for z_M to coalesce with one of z_{B0} or z_{B1} before the outer two roots meet one another? No: the M solution has the same slope at each endpoint, so if it coalesces with either of the other two roots, all three must coalesce simultaneously. Generically, this leads us to expect that in a small neighborhood of the bifurcation point, the error function $y(1) + 1$ plotted in Figure 6 should be approximately a cubic of the form

$$(7.3) \quad y(1) + 1 \approx A(y'(0))^3 + B(\epsilon - \epsilon_c)y'(0).$$

We further expect $A > 0$ and $B > 0$ by examining the shape of the curve in Figure 6. For $\epsilon > \epsilon_c$, there is only one real solution; for $\epsilon < \epsilon_c$, there are three real solutions;

and as $\epsilon \rightarrow \epsilon_c$ from below, these three solutions degenerate into a triple root. All of these expectations are confirmed by the classic pitchfork scenario seen in the lower panel of Figure 8.

7.2. Calculating the Pitchfork Bifurcation Value. The preceding analysis told us that the bifurcation happens when (1.1) admits a solution with $y'(0) = 0$. We can now use that result to determine ϵ_c .

To do so, we first observe that the solution of (1.1) when $\epsilon = \epsilon_c$ must satisfy the boundary condition $y(0) = 1$ as well as the bifurcation condition $y'(0) = 0$. These conditions then uniquely determine the corresponding critical trajectory: it starts at $(y(0), z(0)) = (y(0), y'(0)) = (1, 0)$ and also satisfies $2\epsilon[z + \log(1 - z)] = y^2 - C^2$ by (6.8). So by plugging in $y(0) = 1$ and $z(0) = 0$ we see that $C = 1$ on the critical trajectory.

Next, we use $C = 1$ to find two conditions on ϵ_c . Both conditions relate ϵ_c to z_c , defined as the minimum (i.e., most negative) value of z on the critical trajectory. By solving those two conditions simultaneously, we find z_c and ϵ_c , as follows.

To obtain the first condition on ϵ_c , we note that $z' = 0$ when $z(t)$ reaches its minimum along the trajectory, which implies (from (4.1)) that $y = 0$ there. Like all other points on the critical trajectory, this point $(y, z) = (0, z_c)$ must satisfy (6.8) with $C = 1$. This gives us our first condition on ϵ_c :

$$(7.4) \quad 2\epsilon_c [z_c + \log(1 - z_c)] = -1.$$

The second condition is that the time required for the critical trajectory to go from $y = 1$ to $y = -1$ is $t = 1$, just as it is for every solution of our original boundary-value problem. But for the critical trajectory, we can say more. The critical trajectory is self-symmetric, which implies that the time to go halfway is simply $t = 1/2$. And at that halfway point, the trajectory is at the point we have just been discussing, $(y, z) = (0, z_c)$. To translate these observations into the condition we seek, we need to find a formula for the travel time. The trick is to write dt in terms of z on the trajectory and then integrate. Recall that $z' = y(z - 1)/\epsilon$, from (4.1), so $dt = dz/z' = \epsilon dz/[y(z - 1)]$. Thus

$$\frac{1}{2} = \int_0^{\frac{1}{2}} dt = \int_0^{z_c} \frac{\epsilon_c dz}{y(z - 1)},$$

where we are thinking of y as a function of z on the critical trajectory. That function can be written explicitly by using the fact that $C = 1$ on the critical trajectory; solving $2\epsilon_c[z + \log(1 - z)] = y^2 - 1$ for y and noting that $y \geq 0$ on the first half of the trajectory, we take the positive square root and obtain

$$y = \sqrt{1 + 2\epsilon_c[z + \log(1 - z)]}.$$

Hence our travel time condition becomes

$$(7.5) \quad \frac{1}{2} = \int_{z_c}^0 \frac{\epsilon_c dz}{(1 - z)\sqrt{1 + 2\epsilon_c[z + \log(1 - z)]}}.$$

Now we have two conditions, (7.4) and (7.5), in terms of z_c and ϵ_c . We use (7.4) to eliminate ϵ_c , leaving an integral equation for z_c alone to satisfy. After a bit of algebra, this combined condition is

$$(7.6) \quad 0 = g(z_c) = z_c + \log(1 - z_c) + \int_{z_c}^0 \frac{dz}{(1 - z)\sqrt{1 - \frac{z + \log(1 - z)}{z_c + \log(1 - z_c)}}}.$$

This is now a root-finding problem for the function $g(z_c)$. We find $z_c \approx -3.9052637703$. Plugging this root back into (7.4) yields $\epsilon_c \approx 0.2159869288903$.

7.3. A Deeper Look at the Pitchfork Bifurcation. In retrospect, there were at least two reasons to expect that a pitchfork bifurcation could occur in our problem. The reasons have to do with the symmetry of the problem itself and the special structure of the associated vector field.

First, recall that the differential equation $\epsilon y'' = yy' - y$ and its boundary conditions $y(0) = 1$, $y(1) = -1$ are left unchanged by the transformation $(x, y) \rightarrow (1 - x, -y)$. If we apply this transformation a second time, we get back to (x, y) , as if toggling a switch or reflecting an image in a mirror twice. Such a transformation is known as a \mathbb{Z}_2 symmetry. Now it turns out that the occurrence of a pitchfork bifurcation in a boundary-value problem with \mathbb{Z}_2 symmetry is a codimension-1 phenomenon [20], which means that we should expect to see it in a generic one-parameter family of such problems. So we should not be surprised to find a pitchfork occurring in our problem as we vary its single parameter, ϵ .

Second, recall from section 6.3.1 that the dynamical system (4.1) is Hamiltonian for $z < 1$. McLachlan and Offen [20] give the generic bifurcations for Hamiltonian boundary-value problems and find that a pitchfork can be codimension-1 in planar problems even without the kind of symmetry seen in our system.

So if we had known what we know now, we should have expected a pitchfork all along. As always, everything becomes clearer in hindsight. Fortunately, before clarity comes, we have the pleasure of being surprised.

Appendix A. Higher-Order Asymptotics. Here we collect the long expressions obtained at higher orders of perturbation theory. For the middle-layer solution M , at first order, we find

$$(A.1) \quad Y_1^M = \frac{1}{24} \operatorname{sech}^2 \left(\frac{3X}{4} \right) \left[16 \operatorname{Li}_2 \left(-e^{-3X/2} \right) + 24c_2 + 3X(3X + 4 + 4c_1 - 8 \log 2) \right. \\ \left. + 8 \sinh \left(\frac{3X}{2} \right) \left(2 \log \cosh \left(\frac{3X}{4} \right) - 1 + c_1 \right) \right],$$

where $X \equiv (x - 1/2)/\epsilon$. As discussed in the main text below (5.21), after matching, the integration constants take the values $c_1 = \pi^2/18$ and $c_2 = 1 + \log 4$.

For the B0 case, we condense the notation a bit by defining $X \equiv x/\epsilon$ and $\bar{X} \equiv X - \tanh^{-1} \frac{1}{2}$, so we can write

$$(A.2) \quad Y_1^{B0} = \frac{1}{4} \operatorname{sech}^2 \bar{X} \left[4c_2 + 2\bar{X}(\bar{X} + 1 + c_1 - \log 4) \right. \\ \left. + \sinh(2\bar{X}) \left(c_1 - 1 + 2 \log \cosh(\bar{X}) \right) + 2 \operatorname{Li}_2 \left(-e^{-2\bar{X}} \right) \right].$$

After matching we fix the integration constants as $c_1 = 1 + \log 12$ and

$$(A.3) \quad c_2 = \frac{1}{8} \left(-4 \operatorname{Li}_2(-3) + (\log 3)^2 + \frac{4}{3} \log 6912 \right) \approx 2.594.$$

As discussed in the main text below (5.27), the outer solution and the overlap solution cancel out when we form the composite solution, leaving only the inner solution. Thus

the first-order composite solution for B0 is just

$$(A.4) \quad y_{c,1}^{B0} = -2 \tanh \bar{X} + \epsilon Y_1^{B0},$$

where, as above, $\bar{X} = \frac{x}{\epsilon} - \tanh^{-1} \frac{1}{2}$.

For the B1 solution, we can use the symmetry to write $y^{B1}(x) = -y^{B0}(1-x)$.

Acknowledgments. This paper was written during the global COVID-19 pandemic, which has affected different members of society inequitably. The authors were quite privileged to have had the opportunity to collaborate and continue to do research—opportunities not afforded to everyone. At the same time, one impetus for this work was that S.H.S. taught his course on asymptotics and perturbation methods online, and L.C.S. became aware of Strogatz’s lectures because of Twitter and was able to follow along on YouTube (the problem at hand was discussed in [25]). L.C.S. would therefore like to thank Twitter and YouTube for facilitating his initial involvement in this research. The authors would also like to thank John Boyd, Mathieu Desroches, John Guckenheimer, Martin Krupa, Christian Kuehn, Ian Lizarraga, and Richard Rand for their helpful advice. In their numerical work they made use of the `python` packages `mpmath` [15] for arbitrary precision calculation, `numpy` [10], `scipy` [26], and `matplotlib` [13].

REFERENCES

- [1] See Supplemental Material at <https://github.com/duetosymmetry/surprises-in-a-classic-BVP> for a supplementary `jupyter` notebook demonstrating the numerical approach and a supplementary `Mathematica` notebook demonstrating the analytical calculations. (Cited on pp. 292, 299, 304)
- [2] C. M. BENDER AND S. A. ORSZAG, *Advanced Mathematical Methods for Scientists and Engineers: Asymptotic Methods and Perturbation Theory*, Springer, 1999. (Cited on pp. 291, 292)
- [3] E. BENOIT, J. L. CALLOT, F. DIENER, AND M. DIENER, *Chasse au canard*, *Collect. Math.*, 31 (1981), pp. 37–119. (Cited on p. 292)
- [4] J. BOYD, *Weakly Nonlocal Solitary Waves and Beyond-All-Orders Asymptotics: Generalized Solitons and Hyperasymptotic Perturbation Theory*, Kluwer, 1998. (Cited on p. 293)
- [5] J. P. BOYD, *The devil’s invention: Asymptotic, superasymptotic and hyperasymptotic series*, *Acta Appl. Math.*, 56 (1999), pp. 1–98. (Cited on p. 293)
- [6] R. M. CORLESS, G. H. GONNET, D. E. G. HARE, D. J. JEFFREY, AND D. E. KNUTH, *On the Lambert W function*, *Adv. Comput. Math.*, 5 (1996), pp. 329–359, <https://doi.org/10.1007/BF02124750>. (Cited on pp. 307, 308)
- [7] M. DIENER, *The canard unchained or how fast/slow dynamical systems bifurcate*, *Math. Intelligencer*, 6 (1984), pp. 38–49, <https://doi.org/10.1007/BF03024127>. (Cited on p. 292)
- [8] F. DUMORTIER, R. ROUSSARIE, AND R. H. ROUSSARIE, *Canard Cycles and Center Manifolds*, *Mem. Amer. Math. Soc.* 577, AMS, 1996. (Cited on p. 292)
- [9] W. ECKHAUS, *Relaxation oscillations including a standard chase on French ducks*, in *Asymptotic Analysis II*, *Lecture Notes in Math.* 985, Springer, 1983, pp. 449–497, <https://doi.org/10.1007/BFb0062381>. (Cited on p. 292)
- [10] C. R. HARRIS, K. J. MILLMAN, S. J. VAN DER WALT, R. GOMMERS, P. VIRTANEN, D. COUNAPEAU, E. WIESER, J. TAYLOR, S. BERG, N. J. SMITH, R. KERN, M. PICUS, S. HOYER, M. H. VAN KERKWIJK, M. BRETT, A. HALDANE, J. F. DEL RÍO, M. WIEBE, P. PETERSON, P. GÉRARD-MARCHANT, K. SHEPPARD, T. REDDY, W. WECKESSER, H. ABBASI, C. GOHLKE, AND T. E. OLIPHANT, *Array programming with NumPy*, *Nature*, 585 (2020), pp. 357–362, <https://doi.org/10.1038/s41586-020-2649-2>. (Cited on p. 314)
- [11] E. J. HINCH, *Perturbation Methods*, *Cambridge Texts Appl. Math.*, Cambridge University Press, 1991, <https://doi.org/10.1017/CBO9781139172189>. (Cited on p. 291)
- [12] M. H. HOLMES, *Introduction to Perturbation Methods*, *Texts Appl. Math.*, Springer, New York, 1995. (Cited on pp. 291, 292, 294, 299)
- [13] J. D. HUNTER, *Matplotlib: A 2D graphics environment*, *Comput. Sci. Engrg.*, 9 (2007), pp. 90–95, <https://doi.org/10.1109/MCSE.2007.55>. (Cited on p. 314)

- [14] H. JARDÓN-KOJAKHMETOV AND C. KUEHN, *On fast–slow consensus networks with a dynamic weight*, *J. Nonlinear Sci.*, 30 (2020), pp. 2737–2786, <https://doi.org/10.1007/s00332-020-09634-9>. (Cited on p. 292)
- [15] F. JOHANSSON ET AL., *mpmath: A Python Library for Arbitrary-Precision Floating-Point Arithmetic*, Version 1.2.1, 2013, <https://mpmath.org/>. (Cited on p. 314)
- [16] J. K. KEVORKIAN AND J. D. COLE, *Multiple Scale and Singular Perturbation Methods*, *Appl. Math. Sci.* 114, Springer Science & Business Media, 2012. (Cited on pp. 291, 292)
- [17] M. KRUPA AND P. SZMOLYAN, *Extending geometric singular perturbation theory to nonhyperbolic points—fold and canard points in two dimensions*, *SIAM J. Math. Anal.*, 33 (2001), pp. 286–314, <https://doi.org/10.1137/S0036141099360919>. (Cited on p. 292)
- [18] C. KUEHN, *Multiple Time Scale Dynamics*, *Appl. Math. Sci.* 191, Springer, 2015. (Cited on pp. 291, 292)
- [19] P. A. LAGERSTROM, *Matched Asymptotic Expansions: Ideas and Techniques*, *Appl. Math. Sci.* 76, Springer Science & Business Media, 2013. (Cited on pp. 291, 292)
- [20] R. I. MCLACHLAN AND C. OFFEN, *Bifurcation of solutions to Hamiltonian boundary value problems*, *Nonlinearity*, 31 (2018), pp. 2895–2927, <https://doi.org/10.1088/1361-6544/aab630>, <https://arxiv.org/abs/1710.09991>. (Cited on p. 313)
- [21] F. W. J. OLVER, A. B. OLDE DAALHUIS, D. W. LOZIER, B. I. SCHNEIDER, R. F. BOISVERT, C. W. CLARK, B. R. MILLER, B. V. SAUNDERS, H. S. COHL, AND M. A. McCLAIN, EDs., *NIST Digital Library of Mathematical Functions*, Release 1.1.2, 2021, <http://dlmf.nist.gov/>. (Cited on p. 303)
- [22] R. E. O'MALLEY, *Singular Perturbation Methods for Ordinary Differential Equations*, *Appl. Math. Sci.* 89, Springer, 1991. (Cited on p. 291)
- [23] W. H. PRESS, S. A. TEUKOLSKY, W. T. VETTERLING, AND B. P. FLANNERY, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed., Cambridge University Press, New York, 2007. (Cited on p. 304)
- [24] S. H. STROGATZ, *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*, CRC Press, 2015. (Cited on pp. 292, 296)
- [25] S. H. STROGATZ, *Lecture 16: A Tricky Nonlinear Boundary-Value Problem*, https://www.youtube.com/watch?v=uRhCeg_SR9I, April 2021. (Cited on p. 314)
- [26] P. VIRTANEN, R. GOMMERS, T. E. OLIPHANT, M. HABERLAND, T. REDDY, D. COURNAPEAU, E. BUROVSKI, P. PETERSON, W. WECKESSER, J. BRIGHT, S. J. VAN DER WALT, M. BRETT, J. WILSON, K. J. MILLMAN, N. MAYOROV, A. R. J. NELSON, E. JONES, R. KERN, E. LARSON, C. J. CAREY, Í. POLAT, Y. FENG, E. W. MOORE, J. VANDERPLAS, D. LAXALDE, J. PERKTOLD, R. CIMRMAN, I. HENRIKSEN, E. A. QUINTERO, C. R. HARRIS, A. M. ARCHIBALD, A. H. RIBEIRO, F. PEDREGOSA, P. VAN MULBREGT, AND SciPY 1.0 CONTRIBUTORS, *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python*, *Nature Methods*, 17 (2020), pp. 261–272, <https://doi.org/10.1038/s41592-019-0686-2>. (Cited on p. 314)