Physical Models of Living Systems

Philip Nelson

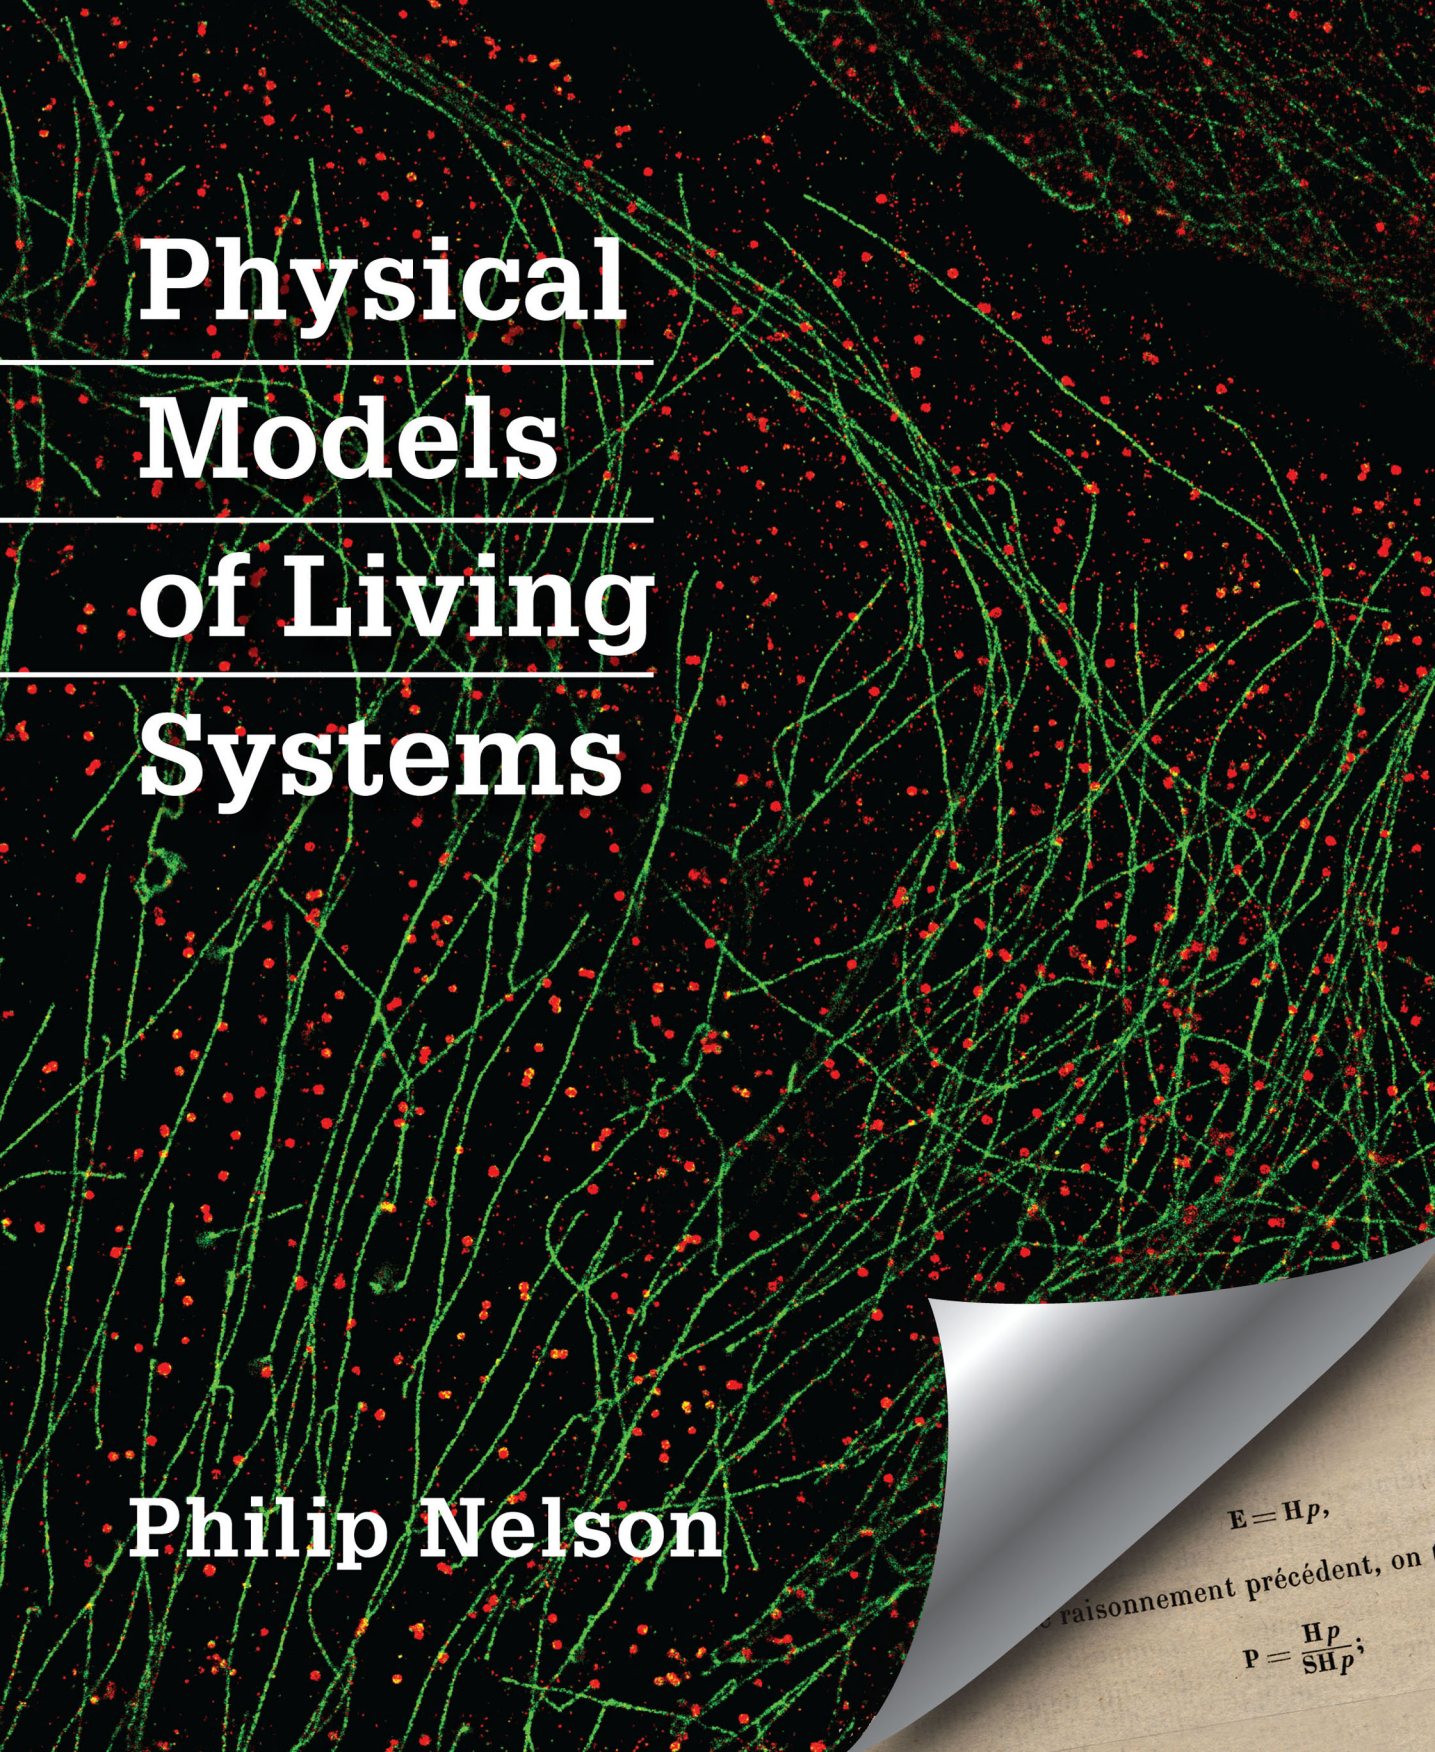

Physical Models of Living Systems

Philip Nelson

University of Pennsylvania

with the assistance of Sarina Bromberg, Ann Hermundstad, and Jason Prentice

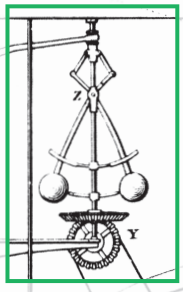

W. H. Freeman & Company

A Macmillan Education Imprint

Publisher: Kate Parker
Acquisitions Editor: Alicia Brady
Senior Development Editor: Blythe Robbins
Assistant Editor: Courtney Lyons
Editorial Assistant: Nandini Ahuja
Marketing Manager: Taryn Burns
Senior Media and Supplements Editor: Amy Thorne
Director of Editing, Design, and Media Production: Tracey Kuehn
Managing Editor: Lisa Kinne
Project Editor: Kerry O'Shaughnessy
Production Manager: Susan Wein
Design Manager and Cover Designer: Vicki Tomaselli
Illustration Coordinator: Matt McAdams
Photo Editors: Christine Buese, Richard Fox
Composition: codeMantra
Printing and Binding: RR Donnelley

Cover: [Two-color, superresolution optical micrograph.] Two specific structures in a mammalian cell have been tagged with fluorescent molecules via immunostaining: microtubules (false-colored *green*) and clathrin-coated pits, cellular structures used for receptor-mediated endocytosis (false-colored *red*). See also Figure 6.5 (page 138). The magnification is such that the height of the letter "o" in the title corresponds to about 1.4 μm. [Image courtesy Mark Bates, Dept. of NanoBiophotonics, Max Planck Institute for Biophysical Chemistry, published in Bates et al., 2007. Reprinted with permission from AAAS.] *Inset:* The equation known today as the "Bayes formula" first appeared in recognizable form around 1812, in the work of Pierre Simon de Laplace. In our notation, the formula appears as Equation 3.17 (page 52) with Equation 3.18. (The letter "S" in Laplace's original formulation is an obsolete notation for sum, now written as $\sum$.) This formula forms the basis of statistical inference, including that used in superresolution microscopy.
Title page: Illustration from James Watt's patent application. The green box encloses a centrifugal governor. [From *A treatise on the steam engine: Historical, practical, and descriptive* (1827) by John Farey.]

Library of Congress Preassigned Control Number: 2014949574
ISBN-13: 978-1-4641-4029-7
ISBN-10: 1-4641-4029-4

©2015 by Philip C. Nelson
All rights reserved

Printed in the United States of America

First printing

W. H. Freeman and Company, 41 Madison Avenue, New York, NY 10010
Houndmills, Basingstoke RG21 6XS, England
www.whfreeman.com

Prolog: A Breakthrough on HIV

Los Alamos, 1994

Alan Perelson was frustrated. For some years, he, and many other researchers, had been staring at an enigmatic graph (Figure 0.1). Like any graph, it consisted of dry, unemotional squiggles. But like any graph, it also told a story.

The enigmatic feature of the graph was precisely what made HIV so dangerous: After a brief spike, the concentration of virus particles in the blood fell to a low, steady level. Thus, after a short, flu-like episode, the typical patient had no serious symptoms, but remained

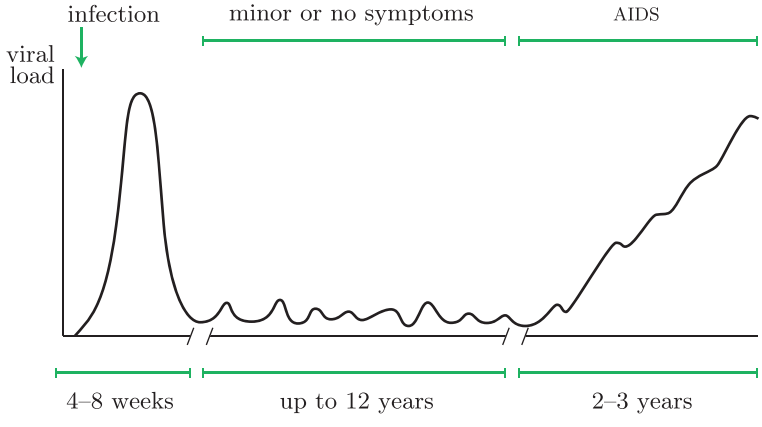

Figure 0.1 [Sketch graph.] **The time course of HIV infection,** representing the progression of the disease as it was understood in the early 1990s. After a brief, sharp peak, the concentration of virus particles in the blood ("viral load") settled down to a low, nearly steady level for up to ten years. During this period, the patient showed no symptoms. Ultimately, however, the viral load increased and the symptoms of full AIDS appeared. [After Weiss, 1993.]

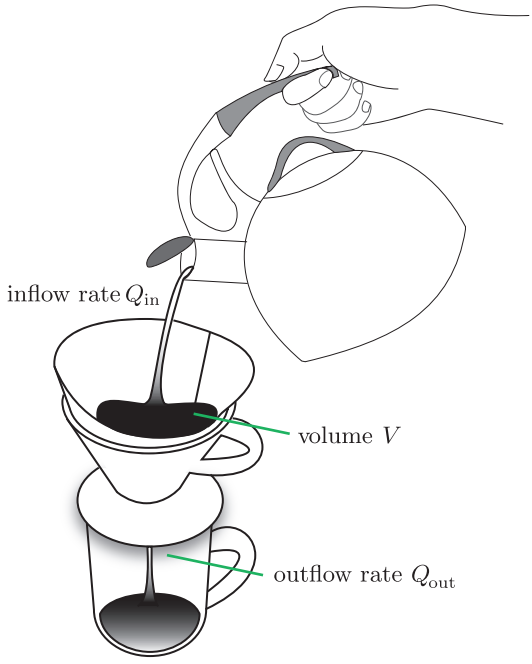

Figure 0.2 [Metaphor.] **Steady state in a leaky container.** Inflow at a rate Q_{in} replenishes the container, compensating outflow at a rate Q_{out}. If we observe that the volume V of liquid in the container is steady, we can conclude that Q_{out} matches Q_{in}, but we can't determine the actual value of either quantity without more information. In the analogy to viral dynamics, Q_{in} corresponds to the body's production of virus particles and Q_{out} to the immune system's rate of virus clearance (see Chapter 1).

contagious, for up to ten years. Inevitably, however, the virus level eventually rose again, and the patient died.

In the early 1990s, many researchers believed that these facts implied that HIV was a slow virus, which remained in the body, nearly dormant, for years before rising sharply in number. But how could such a long latency period be possible? What was happening during those ten years? How could the patient's immune system fight the virus effectively at first, and then ultimately succumb?

Perelson and others had suspected for some time that maybe HIV was not slow or dormant at all during the apparent latent period. He made an analogy to a physical system: If we see a leaky container that nevertheless retains water at some constant level, we can conclude that there must be water flowing into it (Figure 0.2). But we can't determine *how fast* water is flowing in. All we can say is that the rate of inflow equals the rate of outflow. Both of those rates could be small—or both could be large. Applying this idea to HIV, Perelson realized that, during the long period of low blood concentration, the virus might actually be multiplying rapidly, but after the brief initial episode, it could be eliminated by the body just as rapidly.

A real leaky container has another simple property reminiscent of the HIV data: Because the outflow rate $Q_{\text{out}}(V)$ increases as the volume of the water (and hence its pressure at the exit point) goes up, the system can *self-adjust* to a steady state, no matter what inflow rate Q_{in} we select. Similarly, different HIV-infected patients have quite different steady levels of virus concentration, but all maintain that steady level for long periods.

Perelson was head of the Theoretical Biology and Biophysics Group at Los Alamos National Laboratory. By 1994, he had already developed a number of elaborate mathematical models in an attempt to see if they could describe clinical reality. But his models were full of unknown parameters. The available data (Figure 0.1) didn't help very much. How could he make progress without some better knowledge of the underlying cellular events giving rise to the aggregate behavior?

New York City, 1994

David Ho was puzzled. As the head of the Aaron Diamond AIDS Research Center, he had the resources to conduct clinical trials. He also had access to the latest anti-HIV drugs and had begun tests with ritonavir, a "protease inhibitor" designed to stop the replication of the HIV virus.

Something strange was beginning to emerge from these trials: The effect of treatment with ritonavir seemed to be a very *sudden* drop in the patient's total number of virus particles. This was a paradoxical result, because it was known that ritonavir by itself didn't destroy existing virus particles, but simply stopped the creation of new ones. If HIV were really a slow virus, as many believed, wouldn't it also *stay around* for a long time, even once its replication was stopped? What was going on?

Also, it had been known for some time that patients treated with antiviral drugs got much better, but only temporarily. After a few months, ritonavir and other such drugs always lost their effectiveness. Some radically new viewpoint was needed.

Hilton Head Island, 1994

Perelson didn't know about the new drugs; he just knew he needed quantitative data. At a conference on HIV, he heard a talk by one of Ho's colleagues, R. Koup, on a different topic. Intrigued, he later phoned to discuss Koup's work. The conversation turned to the surprising results just starting to emerge with ritonavir. Koup said that the group was looking for a collaborator to help make sense of the strange data they had been getting. Was Perelson interested? He was.

Ho and his colleagues suspected that simply measuring viral populations before and after a month of treatment (the usual practice at the time) was not showing enough detail. The crucial measurement would be one that examined an asymptomatic patient, not one with full AIDS, and that monitored the blood virus concentration *every day* after administering the drug.

More clinical trials followed. Measurements from patient after patient told the same story (Figure 0.3): *Shutting down the replication of virus particles brought a hundredfold drop in their population in 2–3 weeks.*

Perelson and Ho were stunned. The rapid drop implied that the body was constantly clearing the virus at a tremendous rate; in the language of Figure 0.2, Q_{out} was huge. That could only mean that, without the drug, the production rate Q_{in} was also huge. Similar results were soon obtained with several other types of antiviral drugs. The virus wasn't dormant at all; it was replicating like mad. Analysis of the data yielded a numerical value for Q_{out}, as we'll see in Chapter 1. Using this measurement, the researchers estimated that the typical asymptomatic patient's body was actually making at least *a billion* new virus particles each day.[3]

As often happens, elsewhere another research group, led by George Shaw, independently pursued a similar program. This group, too, contained an "outsider" to AIDS

[3]Later, more refined estimates showed that the average production rate was actually even larger than this initial lower bound.

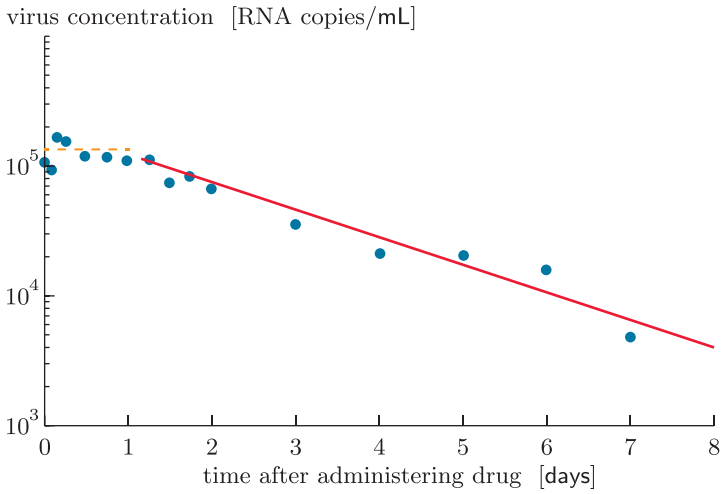

Figure 0.3 [Experimental data with preliminary fit.] **Virus concentration in a patient's blood ("viral load") after treatment with a protease inhibitor**, showing the rapid decline after treatment. In this semilog plot, the *solid line* shows the time course corresponding to elimination of half the total viral population every 1.4 days. The *dashed line* highlights a deviation from this behavior at early times (the "initial plateau"); see Chapter 1. [Data from Perelson, 2002; see Dataset 1.]

research, a mathematician named Martin Nowak. Both groups published their findings simultaneously in *Nature*. The implications of this work were profound. Because the virus is replicating so rapidly, it can easily mutate to find a form resistant to any given drug.[4] Indeed, as we'll see later, the virus mutates often enough to generate every possible single-base mutation every few hours. Hence, every infected patient *already* has some resistant mutant viruses before the drug is even administered; in a couple of weeks, this strain takes over and the patient is sick again. The same observation also goes to the heart of HIV's ability to evade total destruction by the body: It is constantly, furiously, playing cat-and-mouse with the patient's immune system.

But what if we simultaneously administer *two* antiviral drugs? It's not so easy for a virus to sample every possible *pair* of mutations, and harder still to get three or more. And in fact, subsequent work showed that "cocktails" of three different drugs can halt the progression of HIV infection, apparently indefinitely. The patients taking these drugs have not been cured; they still carry low levels of the virus. But they are alive, thanks to the treatment.

The message

This book is about basic science. It's not about AIDS, nor indeed is it directly about medicine at all. But the story just recounted has some important lessons.

The two research groups mentioned above made significant progress against a terrible disease. They did this by following some general steps:

1. Assemble (or join) an interdisciplinary team to look at the problem with different sets of tools;
2. Apply simple physical metaphors (the leaky container of water) and the corresponding disciplines (dynamical systems theory, an area of physics) to make a hypothesis; and

[4] Actually the *fact* of mutation had already been established a few years earlier. Prior to the experiments described here, however, it was difficult to understand how mutation could lead to fast evolution.

3. Perform experiments specifically designed to give new, quantitative data to support or refute the hypothesis.

This strategy will continue to yield important results in the future.

The rest of the book will get a bit dry in places. There will be many abstract ideas. But abstract ideas do matter when you understand them well enough to find their concrete applications. In fact, sometimes their abstractness just reflects the fact that they are so widely applicable: Good ideas can jump like wildfires from one discipline to another. Let's get started.