from *SIAM News*, Volume 32, Number 2

Will Viruses Succumb To Mathematical Models?

By Barry A. Cipra

The things that we have said are not to be looked upon as proofs but rather as indications of the dullness of those who think differently.

—Galen, *On the Natural Faculties*

The early physicians had their theories. Even the great Galen, the father of pharmaceuticals, couldn't resist the urge to philosophize. He expounded a four-humor, three-spirit theory of physiology that held sway for a millennium and a half. Blood, phlegm, and biles both black and yellow were thought to mix it up with the natural, vital, and animal spirits residing in the liver, heart, and brain.

Modern medicine has done away with most of the sophistry that used to pass for scientific reasoning. One of the up-and-coming tools of medical research, in fact, is mathematical modeling: the use of equations to explore the logical consequences of hypotheses and data. Conclusions reached mathematically engender far more confidence than speculations rooted in philosophy—even when the underlying assumptions are known to be wrong!

A case in point is recent work in modeling the long, supposedly quiescent phase characteristic of infections with two viruses: HIV and hepatitis C. A team of researchers led by Alan Perelson, a mathematical immunologist at Los Alamos National Laboratory, has developed mathematical models that account for several striking features in patients' responses to drugs. The team's analyses, coupled with new "high-resolution" experiments, have caused clinicians to rethink their standards for drug therapies.

"Mathematical models allow us to gain insights into processes of viral growth and destruction of host cells that were not readily available before we introduced modeling," Perelson says. In particular, the models show considerable activity in an infected person, even when there are no outward signs of change. That fact alone (properly quantified) has implications for treatment strategies, Perelson points out.

Viral Primer

Asclepiades goes far astray in bidding us distrust our senses where obvious facts plainly overturn his hypotheses. Much better would it have been for him not to assail obvious facts, but rather to devote himself entirely to these.

—Galen, *On the Natural Faculties*

The human immunodeficiency virus (HIV) and hepatitis-C virus (HCV) are blood-borne viruses, most commonly transmitted through sexual contact or needle sharing by intravenous drug users. In both cases, infected patients typically experience a relatively brief "acute" phase, with flu-like symptoms lasting a few weeks to a few months, followed by a long symptom-free period, often lasting a decade or more, during which the viral "load" stays more or less constant, at a level much lower than its initial peak. For most patients infected with HIV, the viral load ultimately increases dramatically again and symptoms of acquired immunodeficiency syndrome appear. (The relationship of HIV to AIDS was initially controversial, but most experts now agree that HIV is the "cause" of AIDS. HIV's favorite "target" is a component of the immune system known as the $CD4^+$ T cell.)

In the case of HCV, cirrhosis develops in 20–30% of infected people, and cancer of the liver in a far smaller percentage. (Results of a 1989 study in Jefferson County, Alabama, suggest that HCV and alcohol abuse are equally important as causes of liver disease.) Hepatitis C is distinct from hepatitis A and B. HAV, also known as infectious hepatitis, is usually eradicated within a few months by the body's own immune system, much as measles is; approximately one-third of the U.S. population shows signs of past infection—i.e., immunity. HBV is more serious. A small percentage of infected adults, but a large percentage of children, become chronically infected, with the attendant risks of liver disease later on. HCV is chronic in nearly all cases.

The genetic material of both HIV and HCV is single-stranded RNA, rather than the double-stranded DNA usually found in living organisms. HIV differs from HCV, however, in being a *retrovirus*: To replicate in a host cell, it first converts its genome to DNA, using an enzyme called *reverse transcriptase*. Another enzyme,*integrase*, then inserts the newly formed DNA into the host cell's genome. In effect, the host cell instantly evolves into a "higher" form of life. These extra steps have theurapeutic implications: Blocking either protein would render the virus noninfectious.

HIV also carries an enzyme of a third type, called a *protease*, which plays an essential role in creating new viral particles. The HIV portion of the host cell's genome, when activated, not only mass-reproduces its original RNA, it also sets up a protein factory. Most of the proteins begin as one big "polyprotein." HIV's protease cuts itself out of the polyprotein and then makes additional cuts in specific places, freeing the individual pieces to do their various jobs, assembling the new virus particles and preparing for a new round of infections. Protease is thus another potential Achilles heel for HIV, since uncut polyprotein is incapable of doing anything.

Current drug therapies for HIV focus on inhibiting either reverse transcriptase or protease. Zidovudine (AZT), one of the earliest drugs for HIV, is a reverse transcriptase inhibitor; ritonavir is one example of a protease inhibitor. Current guidelines recommend treating patients with a "cocktail" of two or three drugs with different targets or mechanisms.

HCV also relies on a protease to replicate itself. (Not being a retrovirus, it doesn't convert its genome to DNA, making that whole

avenue of HIV treatment unavailable for HCV. Viruses in general, however, seem to have latched onto polyproteins and proteases as convenient means for data compression: The alternative is for the genome to include separate "start" and "stop" instructions for each protein. In fact, proteases are ubiquitous in cells as well; runaway protease production has been implicated in such diseases as osteoporosis and asthma.) A protease inhibitor for HCV is under development, but the only drug currently approved for HCV infection is interferon (more precisely, interferon-α-2b, a naturally occurring protein manufactured during the body's own immune response), alone or in combination with ribavirin, another antiviral agent. The standard treatment has been 3 million "international units" (3 mIU) of interferon administered by injection three times a week for a full year. However, for patients who don't respond within 90 days—and most don't—the recommendation has been to discontinue interferon treatment and hope that something better will come along.

Treatment strategies for HIV and HCV remain as much art as science. Despite the great strides in biochemistry, many of the details of each illness remain unclear. For example, it's not known exactly how interferon works. Nor is it known why one patient responds to treatment and anoth-er doesn't. To a large extent, physicians must rely on experience, falling back on the model made famous by Hippocrates: careful observation of the course of the disease and recording of the outcomes of treatments.

A Model Disease

Some of his simples he found useful against aconite, others against sea-hare, and so forth. So Mithradates mixed these all together and made a single drug, hoping he would now have a resource against all poisons.

—*Galen*, On Antidotes

That's where modeling comes in.

Clinical drug trials typically focus on a single, key question: How well does the drug in question work? Perelson and col-leagues have shown that data gathered for that purpose can also be used to ad-dress another fundamental question: What is the virus up to? In particular, measurements of viral load taken during clinical trials shed light on the "natural" rates of viral production and clearance. The numbers are surprising.

In abstract terms, injecting (or swallow-ing) a drug amounts to perturbing a dynamical system. If you have a general notion of the equations that describe the system, studying its response to a perturbation can help pinpoint values of the parameters. The quantitative model can then be used to suggest the possible effects of other perturbations. In short, it provides guidelines for future research.

The virus modelers began studying HIV in 1994. Their first report, "Rapid Turnover of Plasma Virions and CD4 Lymphocytes in HIV-1 Infection," by David Ho, Wen Chen, and Martin Markowitz of the Aaron Diamond AIDS Research Center in New York, Perelson and Avidan Neumann of Los Alamos, and John Leonard of Abbott Laboratories in Illinois, appeared in *Nature* (January 12, 1995). (Researchers are currently concerned with two types of HIV: HIV-1, the more prevalent of the two and the one that has received the most attention, and HIV-2, which is of increasing concern.) A follow-up report, "HIV-1 Dynamics in Vivo: Virion Clearance Rate, Infected Cell Life-span, and Viral Generation Time," appeared in *Science* (March 15, 1996).

The basic model considers the rate of change of three variables: the number of uninfected target cells T, the number of productively infected target cells T^*, and the number of virus particles V. What's measurable experimentally is the total number of target cells, $t + t^*$, and virus particles, v, in a blood sample, usually expressed as a count per unit volume. The model assumes that the lower-case variables are a constant fraction of the upper-case variables. (The model could work with just the measurable variables, but one of the interesting conclusions is based on the order of magnitude of V.)

The underlying equations are relatively simple:

$$dT/dt = s + pT(1 - T/T_{\max}) - d_T T - kVT$$

$$dT^*/dt = kVT - \delta T^*$$

$$dV/dt = N\delta T^* - cV$$

In other words, in the absence of virus, the T cells obey a standard lo-gistic equation, but with virus pre-sent they become infected at a rate k. Once the infection has set in, the virus kills host T cells at a rate δ, producing N new virus particles per infected cell in the process. Final-ly, the virus itself is cleared at rate c.

After the initial infection, the viral load in most HIV patients reaches a set point and stays essentially constant for years. That implies a balance between viral production and clearance: $N\delta T_0^* = cV_0$, with the subscripts indicating the steady-state values. It's an easy step from there to $NkT_0 = c$. But are the rates large or small?

One way to find out is to perturb the system.

"What modeling allowed us to do was go in and look at what was going on in patients when it appeared that nothing was happening," Perelson explains. The various rate constants can be estimated from patients' responses to drugs.

The modelers analyzed two scenarios, corresponding to treatment with a reverse transcriptase inhibitor and treatment with a protease inhibitor. In the former case, the virus can be viewed as immediately rendered less, or even non-, infective: The constant k is reduced by a factor $1 - \eta_{RT}$, with η_{RT} representing the effectiveness of the reverse transcriptase inhibitor. If the drug is 100% effective, the crucial equations become

$$dT^*/dt = -\delta T^*$$

$$dV/dt = N\delta T^* - cV,$$

which are easy to solve analytically. The viral load decays exponentially: $V(t) = V_0(ce^{-\delta t} - \delta e^{-ct})/(c - \delta)$. The general case is also easy to solve analytically if the T-cell count is assumed to stay at its steady-state value $T_0 = c/Nk$. The viral load again decays exponentially. The assumption of a constant T-cell count, however, is contradicted both by the model and by experimental data: The reality is that T-cell counts tend to increase significantly. Perversely, this ostensible sign of health may actually help the virus survive, albeit at lower levels. A more detailed analysis shows that for the virus to be eliminated, the effectiveness η_{RT} must exceed some critical value.

The scenario for the protease inhibitor is more complicated. Whereas a 100% effective reverse transcriptase inhibitor immediately prevents the infection of new cells, causing an immediate exponential decay in T^*, a protease inhibitor allows the existing, "healthy" virus particles to infect T cells, but their progeny, containing uncleaved polyproteins, will no longer be infectious. In other words, V now consists of two terms: an exponentially declining infectious term, $V_I(t) = V_0 e^{-ct}$, and a noninfectious term governed by the equation $dV_{NI}/dt = N\delta T^* - cV_{NI}$. The equation for infect-ed T cells is $dT^*/dt = kV_I T - \delta T^*$. If T is again assumed to be constant—which is reasonable for a short period of time—the total viral load is found to be

$$V(t) = V_0 e^{-ct} + \frac{cV_0}{c-\delta}\left[\frac{c}{c-\delta}\left(e^{-\delta t} - e^{-ct}\right) - \delta t e^{-ct}\right].$$

In the paper that appeared in *Science*, the modelers analyzed data from five HIV-infected patients who were receiving the protease inhibitor ritonavir. Blood samples were taken every few hours for the first two days and then daily for a week (see Figure 1). Based on their anal-ysis, the modelers concluded that the half-life of virus par-ticles is about five hours, and that of in-fected T cells about a day and a half.

Possibly most significant are the estimates for viral production. The researchers concluded that prior to treatment, pa-tients are producing (and clearing) on the order of 10 billion virus particles per day.

The significance of this number lies in its implications for drug resistance. Experiments with cell cultures indicate that when HIV replicates, errors occur at a frequency of 3×10^{-5} per base per generation. The HIV genome is about 10,000 bases long, and thus on average there is a 30% chance that one base will be changed. Since each base can be changed to one of three others (there are only four bases in RNA or DNA), this means that there are approximately 30,000 possible single-error mutations of a given genome. But tens or hundreds of millions of cells must be newly infected in order to pro-duce the 10 billion viruses each d This suggests that every single-error v ant is produced in large quantities o daily basis. Consequently, a drug tha designed to bind with a particular am acid at a certain location in a prot can't be counted on to work for v long.

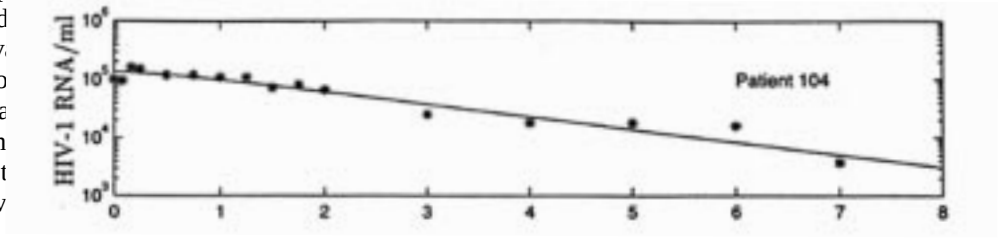

Figure 1. *Viral load in an HIV patient (*) compared with the predictions of a mathematical model. (Figure courtesy of Alan Perelson.)*

"[The model] gave us a real, rational understanding, for the first time, of why the drugs being used to treat AIDS patients were not very effective," Perelson says. "Not only are all [single-error] mutations going to be explored, they're going to be explored in reasonably large numbers." The findings hint at a quantitative basis for designing resistance-resistant treatment schemes. "The idea is if you use at least three drugs, you're typically building a barrier where you might need five or more mutations before you would generate high levels of resistance," Perelson explains.

The findings also suggest that immedi-ate treatment, when the viral population is still very homogeneous, is preferable to a wait-and-see approach. However, given the drawbacks of HIV therapies—these extremely potent drugs have serious side effects—the potential advantage of early treatment is far from clear-cut. Clinical trials are still the final arbiter.

Rethinking Hepatitis

Do we still, then, disbelieve that each drug attracts that humour which is proper to it?

—Galen, *On the Natural Faculties*

The model for hepatitis C is similar to the HIV model. In a paper published in *Science* (October 2, 1998), "Hepatitus C Viral Dynamics in Vivo and the Antiviral Efficacy of Interferon-α Therapy," Perelson and colleagues Neumann, Harel Dahari of Bar-Ilan University in Israel, David Gretch of the University of Washington, and Nancy Lam, Thomas Lay-den, and Thelma Wiley of the University of Illinois at Chicago analyzed data from a study of 23 patients infected with HCV.

The patients were divided into three groups and given daily injections of 5, 10, or 15 mIU of interferon for two weeks, after which all three groups received 5 mIU daily. (The current standard treatment is 3 mIU.) Blood samples were taken every few hours for the first two days and then daily for two weeks.

The data show a decidedly biphasic decline in virus: a precipitous drop in the first day or so, and a slower decline there-after.

Moreover, the slopes of the decline curve clearly show dose-dependence.

The modelers conclude from these results that interferon acts by blocking the production or release of new virus particles by existing infected cells, rather than by preventing the infection of new cells or increasing the death rate of infected cells. Their analysis indicates a blocking efficacy of 81, 95, and 96% for the 5-, 10-, and 15-mIU regimes, respectively, and a virus half-life of approximately 2.7 hours. (The estimated half-life for infected cells ranged from two days to two months, correlated with the initial amount of virus. The viral half-life estimates ranged from 1.5 to 4.6 hours.) The model also indicates a staggering amount of (pretreatment) viral production and clearance: a *trillion* virus particles per day.

"There are other things of clinical interest in the hepatitis study," Perelson points out. In particular, the slope of the slow, second-phase decline over the two-week period appears to be a useful indicator of how well the treatment is working. "We found that it was predictive of whether or not the patients would 'go undectable' by 90 days," Perelson explains. By current standards, if the viral load hasn't fallen below detectable levels within 90 days of interferon therapy, the treatment is considered unsuccessful and the therapy is stopped. "We're suggesting that you don't even have to wait three months," Perelson says.

That, as usual, remains to be seen. The virus models incorporate numerous simplifying assumptions. Further research will test the limits of their explanatory power. A hundred or a thousand years from now, scientists may find today's theories as fanciful as we find Galen's. But the value of careful thought and study is beyond question. As Galen, himself quoting an earlier author, put it:

People who are unused to learning, learn little, and that slowly, while those more accustomed do much more and do it more easily. The same thing also happens in connection with research. Those who are altogether unfamiliar with this become blinded and bewildered as soon as their minds begin to work: they readily withdraw from the inquiry, in a state of mental fatigue and exhaustion, much like people who attempt to race without having been trained. He, on the other hand, who is accustomed to research, seeks and penetrates everywhere mentally, passing constantly from one topic to another; nor does he ever give up his investigation; he pursues it not merely for a matter of days, but throughout his whole life.

Barry A. Cipra is a mathematician and writer based in Northfield, Minnesota.