



The Mathematics of Mosquitoes and West Nile Virus

Marjorie Wonham[†]

Lying on my sleeping pad, I warily eye the mosquito perched above my head. I could reach up and squash it, but that would require extracting my arm from the warmth of my sleeping bag. So for now, it clings to the yellow nylon of my tent, unaware of its reprieve. Although I may triumph over this particular mosquito, I am all too aware of being vastly outnumbered outside my tent.

Until recently, my interest in mosquitoes was largely pragmatic: avoid, repel, or swat. Lately, though, I have developed a grudging curiosity about how they make a living. Fact: a female mosquito overwinters with fertilized eggs so the first thing she does in spring, before even feeding, is lay eggs. Fact: Even if she doesn't find a blood meal, she can survive by sucking plant juices. Fact: the combined meals of a mosquito horde can (and do) bleed a newborn calf to death. Fact: at best a mosquito bite simply itches; at worst, it means disease transmission—malaria, yellow fever, dengue fever, and West Nile virus are all mosquito-borne.

It is West Nile virus that has piqued my interest of late. First identified in Uganda in 1937, the virus is well established in its native Africa where it lives primarily in birds and is transmitted among them by mosquitoes. Only occasionally does a mosquito transmit the infection to a mammal. From time to time a West Nile virus outbreak occurs in Europe and Africa—in Israel in the 1950s and South Africa in 1974, and more recently in Romania, Morocco, Tunisia, Italy, France, and Russia. Just recently, West Nile virus made its first known, and headline grabbing, North American appearance.

In the summer of 1999, the birds of New York City began mysteriously to die, their bodies appearing conspicuously in the city zoo, parks, and backyards. At first the cause was unknown, but by December of the same year it had been identified, in two reports published in the same issue of Science magazine, as West Nile virus, a disease never before seen on this continent. In subsequent summers, West Nile virus spread west across the continent reaching Ontario in 2001, California and Washington in 2002, and Alberta in 2003.

Corvids—crows and jays—were the hardest hit among the birds; other passerines such as sparrows also carried the virus but were dying in smaller numbers. Among mammals, horses appeared especially vulnerable, with a mortality rate of approximately 40%. Human cases were less common and less likely to be fatal, but were a growing health concern nonetheless. By the end of 2003, the virus had been identified in 7 Canadian provinces and 46 U.S. states, in at least 10

mosquito, 150 bird, and 17 other vertebrate species, and in a total—for that year alone—of over 11,000 human cases in Canada and the U.S.

In Alberta, West Nile virus was first reported in 2003. The year before, I had moved to Edmonton to join a group of mathematical biologists at the University of Alberta. There, I was immersed in a world of mathematical modelling used to tackle biological questions. Two years later, here in my tent on a canoe trip in the Northwest Territories, I have a chance to reflect on the unpredicted collaboration that developed with my mathematical colleagues on the dynamics and control of West Nile virus.

At first glance, we made an unlikely trio for this project. Tomás was a graduate student and programming whiz who studied chamomile invasions on farmland. Mark was a mathematical biologist who had modelled the movement of birds and wolves. And I was a marine ecologist with a mathematical background largely limited to reading tide tables. But together we were galvanized by a question from a colleague in Ontario: “How come no one is modelling West Nile virus?” Hugh asked. And with that casual question began our Year of the Mosquito.

How could mathematical modelling help us understand West Nile virus dynamics? The virus was spreading, and control proposals were beginning to include spraying adult and larval mosquitoes, removing larval mosquito habitat, and even removing birds. Since all of these strategies would cause additional impacts on the environment, perhaps modelling could help maximize control effectiveness while minimizing unwanted effects? I had no idea where to begin, but luckily, I was in good mathematical hands.

Mark and Tomás introduced me to a class of disease models known by their acronym as SIR models (see accompanying articles by Fred Brauer and David Earn). These models were first extended to vector-borne diseases by R. Ross in the early 1900s and G. Macdonald in the 1950s, to combat malaria. Since then, a large associated body of mathematical theory, and an impressive history of contributing to disease control, have both evolved.

In an SIR model, the host population is divided into three groups: Susceptible (healthy uninfected individuals), Infectious (infected and capable of transmitting the disease), and Removed (immune, dead, or otherwise removed individuals). The rates at which an average individual moves from Susceptible to Infectious and from Infectious to Removed are determined, and the relevant birth and death rates are incorporated.

Once constructed, a key piece of information can be extracted from an SIR model, called the disease *reproduction number*, or \mathcal{R}_0 (“R-zero” or “R-naught”). \mathcal{R}_0 tells us the number of new infections that would result from the introduction of a single infectious individual into an entirely susceptible population. For example, if a student with chicken pox walked into a classroom of individuals with no previous exposure to the disease, how many new cases would be caused by direct contact with the initial infectious student? The answer is given by \mathcal{R}_0 . Or for West Nile virus, if an infectious bird arrived in a new city, how many other birds would be infected (via mosquito bites) by that original bird? Again, \mathcal{R}_0 tells us the answer.

The expression for \mathcal{R}_0 is constructed, according to a particular formula, from variables and parameters in the model. Reasonably enough, it takes into account factors such as how long the first individual remains infectious, the likelihood of contact between the infectious and susceptible individuals (either directly or via another species), and how often contact leads to disease transmission. Details of the mathematics underlying the calculation of \mathcal{R}_0 are given in the article by Fred Brauer.

[†] Marjorie Wonham is a Postdoctoral Fellow in the Centre for Mathematical Biology and the Department of Mathematical & Statistical Sciences at the University of Alberta. Her e-mail address is mwonham@ualberta.ca.



BY SWITCHING TERMS IN HER EQUATIONS, MARJORIE FOUND SHE COULD REVERSE THE WEST NILE VIRUS EPIDEMIC.

©Copyright 2004
Gabriela Novakova

For disease control, the value of \mathcal{R}_0 is key. $\mathcal{R}_0 < 1$ means that an infectious individual will, on average, generate fewer than one new infection, so the disease will die out even without control efforts. On the other hand, $\mathcal{R}_0 > 1$ means that an infectious individual will generate more than one new infection, so a disease outbreak will occur. In this case, the disease can be controlled by methods that alter one or more of the model components—such as mortality, contact, or transmission rates—to reduce \mathcal{R}_0 below one.

Armed with this mathematical background, we were ready to develop an SIR model for West Nile virus in North America, calculate \mathcal{R}_0 , and identify how and how much to control the disease to prevent an outbreak. First, we had to define the biological and geographic scope of our efforts. We had already learned that the virus persisted in transmission cycles between mosquitoes (vectors) and birds (reservoir hosts). Although it occasionally spread to other vertebrates (including humans), it seemed not to return to mosquitoes. In other words, although an infection might be deeply significant to the human in question, it would not influence the overall disease dynamics.

This biological fact helped us simplify the mathematical description: by viewing the virus outbreak level in mosquitoes and birds as a proxy for the human infection risk, we could limit our model to only the vector and the reservoir host. For mosquitoes, we focused on the species group that was emerging as the dominant North American vector (*Culex pipiens* spp.). For birds, we focused on the species with the best available infection and mortality data, the American crow (*Corvus brachyrhynchos*). And since the best virus prevalence data were available where West Nile had first appeared, we confined ourselves to modelling the New York City outbreak. Finally, since the disease outbreak at this latitude showed a marked seasonality, appearing in summer and disappearing in winter, we confined the model to a single season from spring through fall.

Soon we had the skeleton of a model. We defined three groups (S , I , and R) for both mosquitoes and birds. From the literature, we obtained estimates of the mosquito biting rate and the transmission probabilities that allowed a mosquito to


infect a bird (which was around 88%) and a bird to infect a mosquito (which was only around 16%). We had recovery and mortality rates for birds, but not for mosquitoes since they didn't seem to be affected by the disease. Since we had birth and death rates for mosquitoes but not for birds, we assumed the birds reproduced once in spring before the model began, and their background (natural) mortality rate would be negligible in the one summer.

To refine the model and assign numerical values to the parameters, we divided up the work according to our expertise. Tomás investigated how to solve and simulate these models on a computer, and Mark explored the mathematical theory underlying this approach. I searched the biological literature for parameter values for mosquitoes, birds, and virus transmission. I didn't envy Tomás and Mark, as my task seemed much the easiest; I was surprised to learn later that they felt the same way about their roles.

Nonetheless, tracking down the biological data took some sleuthing. Today, from the safety of my tent, I can make a guess as to the local mosquito abundance outside: it's very, very high. If I were to stick out a bare arm, it would be covered almost instantly; if I left it out, I could watch an individual probe repeatedly, biting several times before finding her blood meal. (I don't do this often.) As I canoe down the river, I sometimes see a cloud of mosquito larvae rising to the surface and hatching.

For our model, though, we needed observations like this to be quantified. Just how many mosquitoes were out there? How many eggs did one female lay, and how many larvae hatched? How many crows were there? How many mosquito bites per crow in a day? And how contagious was the virus? My quest was something of a scavenger hunt travelling back through biological history, with each paper leading me to an older one. At the same time, new reports about the virus were appearing almost daily in print and online. I divided my time between dusty library shelves and internet listserves.

New biological information, new equations, new parameter estimates, new model analyses, and new simulations surfaced. For several months, we worked to tailor the model as best we could to the biological information. In the interests of tractability, the model had to remain as simple as possible.



WE DO NOT RECOMMEND SWITCHING VARIABLES IN THE EQUATIONS!

©Copyright 2004
Wieslaw Krawcewicz

But the biological complexity seemed almost infinite. At times, I wondered if we were simplifying too much for the model to be informative, while Tomás and Mark wondered if the biology was making the model too complicated to be useful.

In the end, the biology dictated two substantial additions to the model. We learned that mosquitoes could spend quite a long time, up to 14 days, as aquatic larvae that don't bite birds and therefore don't transmit the disease. It can also take quite a long time, perhaps 10–12 days, for an infected mosquito to develop a viral load high enough to transmit the disease back to a bird. These two time periods could add up to almost half a mosquito's lifespan, so they could substantially alter the disease dynamics. We therefore added two new groups to the population: one for larval L and one for exposed E mosquitoes.

With these additions, it seemed we finally had a model we all felt was realistic and tractable. This is the model illustrated in Figure 1. Now that we had the model, it was time for a test: would it behave realistically? We had to know this before we could use \mathcal{R}_0 to make any predictions.

As a test, we chose the records of West Nile virus incidence in both mosquitoes and birds from New York City in 2000. We plugged our literature-based parameter estimates into the model, crossed our fingers, and ran the simulations. Sure

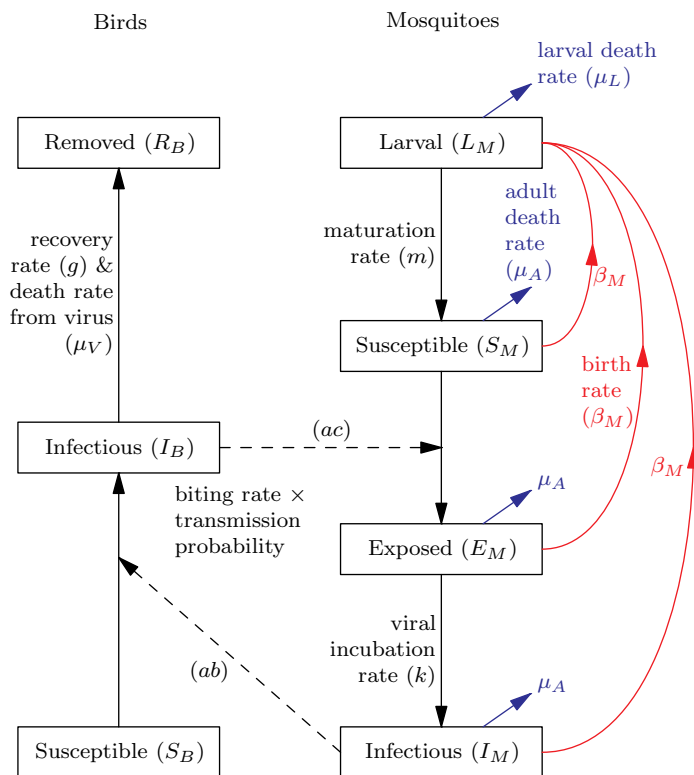


Figure 1: Cartoon illustration of SIR model for West Nile virus cross-infection between birds and mosquitoes. The boxes represent *variables*: the number (or proportion) of birds that are susceptible, infectious, and removed, and mosquitoes that are larval, susceptible, exposed, and infectious. The arrows represent *parameters*: the daily rates at which the value of each variable changes. Solid black lines show the rates (and reasons) for individual birds and mosquitoes moving from one category to another. The two dashed lines represent the rates at which mosquitoes infect birds and birds infect mosquitoes. The red and blue arrows respectively show the birth and death rates for mosquitoes (which can reproduce multiple times during a single season).




Figure 2: Numerical simulations of crow proportions and mosquito numbers predicted by the West Nile virus model for a steady-state mosquito population ($S_{M_0}/N_{B_0} = 30$) over a season lasting 14τ days, where τ is a time unit representing 9.4 days (this time unit may seem funny, but it is a result of manipulating the equations to a nondimensional form—a detail that is not shown here). a) the proportion of susceptible (solid line) and removed birds (dashed line) on the left axis and infectious birds (red line) on the right. b) the number of susceptible (solid line) and larval mosquitoes per bird (dashed line) on the left axis and exposed (dashed red line) and infectious mosquitoes per bird (solid red line) on the right.

enough, for a reasonable range of starting values the model predicted a range of mosquito and bird infection levels that matched the observed data (Figure 2). We were now in a position to use the model to investigate control strategies.

Following methods developed by other mathematicians, we calculated \mathcal{R}_0 for our model, which turned out to be made up of three elements:

$$\mathcal{R}_0 = \frac{ab}{\mu_A} \cdot \frac{ac}{\mu_V + g} \cdot \frac{S_{M_0}}{N_{B_0}} \left(\frac{k}{k + \mu_A} \right).$$

The parts work together like this. The first factor represents the number of new crow infections caused by an infectious mosquito. This is the number of bites per day per mosquito a multiplied by the probability of the virus being transmitted from an infectious mosquito to a crow b multiplied by the average number of days $1/\mu_A$ a mosquito would live.

The second factor is the mirror image: the number of new mosquito infections caused by an infectious crow. The mosquito biting rate a is multiplied by the transmission probability from crows to mosquitoes c and by the average number of days until the infectious crow either recovered ($1/g$) or died ($1/\mu_V$).

The third factor represents, generally speaking, the number of infectious mosquitoes per crow. Specifically, it is the number of initially susceptible mosquitoes S_{M_0} that survive the virus exposure period $k/(k + \mu_A)$ for every bird N_{B_0} . Together, these three elements give us the expression for the total disease \mathcal{R}_0 from birds to birds (via mosquitoes) or from mosquitoes to mosquitoes (via birds). Taking the square root of the right hand side of the expression for \mathcal{R}_0 is a common convention that gives the geometric mean \mathcal{R}_0 from bird to mosquito and *vice versa*.

(If you're looking at Figure 1 and wondering why the mosquito birth rate and larval death rate don't seem to show up in \mathcal{R}_0 , rest assured. They are accounted for by the equality $L_{M_0} = \beta_M S_{M_0}/(m + \mu_L)$, a simplifying assumption in the model that ensures a constant mosquito population.)

Our parameter values gave \mathcal{R}_0 greater than one, predicting a disease outbreak. We were now able to return to our original questions, namely, how could West Nile virus be controlled, and how much control was needed?

One possible answer was obvious: if every mosquito and bird were removed, the virus could not persist. But we were hoping to find a more palatable answer. Examining the expression for \mathcal{R}_0 , we found the ratio S_{M_0}/N_{B_0} in the numerator, telling us that reducing the number of mosquitoes S_{M_0} would reduce \mathcal{R}_0 , but reducing the number of birds N_{B_0} would not. In fact, reducing bird abundance would only make things worse by increasing the value of \mathcal{R}_0 . This was our first lesson: *reducing mosquitoes could help control the virus, but removing birds would only increase the chances of an outbreak.*

Would every mosquito have to be removed to prevent an outbreak? This was our second lesson. Plotting a graph of \mathcal{R}_0 vs. S_{M_0}/N_{B_0} showed us that *the virus could be controlled simply by reducing mosquito abundance, without requiring that every last individual be eliminated* (Figure 3).

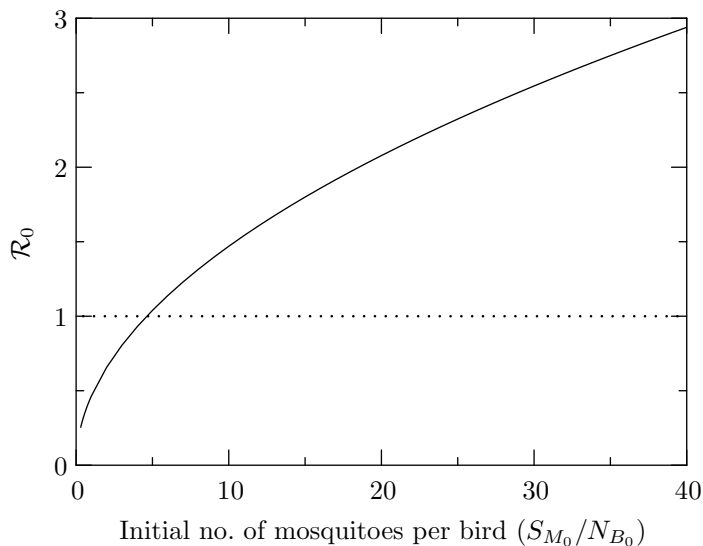


Figure 3: Plot of \mathcal{R}_0 vs. the initial number of mosquitoes per bird S_{M_0}/N_{B_0} , showing that the mosquito population can simply be reduced, and not completely eliminated, to bring \mathcal{R}_0 below one and therefore prevent a West Nile outbreak.

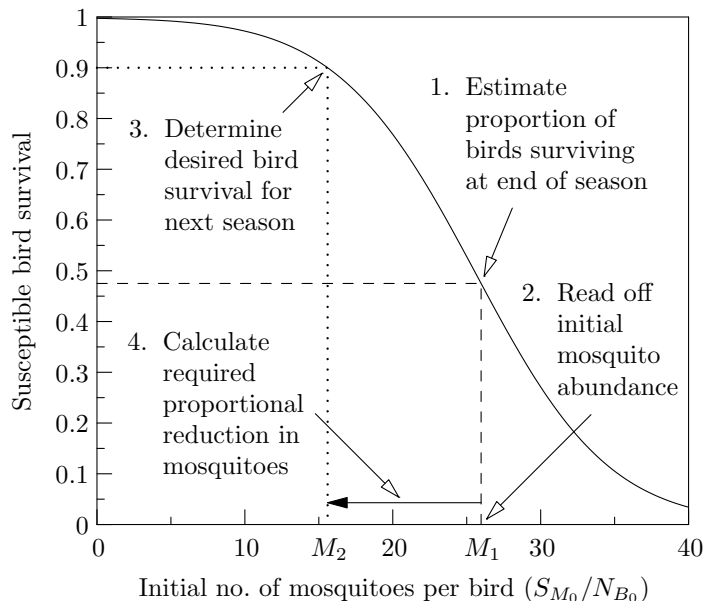


Figure 4: Plot of susceptible bird survival at the end of the season S_B vs. the initial number of mosquitoes per bird at the beginning of the season S_{M_0}/N_{B_0} , showing how future mosquito control efforts can be estimated from previous bird mortality data.

This begged the final question of exactly how much mosquito reduction was enough. It took us a while to work this out, but eventually we came up with an indirect route to the answer. It requires the following four steps, which are illustrated in Figure 4.

1. Obtain from public health authorities an estimate of the proportion of crows that survived the virus in the last season. (These data are becoming increasingly available.)
2. Determine how many mosquitoes per crow M_1 there must have been at the start of the last season to cause the observed bird mortality.
3. Determine the maximum allowable bird mortality for the following season, and the corresponding maximum allowable initial number of mosquitoes per crow M_2 .
4. The ratio of these two mosquito-per-crow estimates, M_2/M_1 , gives the proportional reduction in mosquitoes needed to control the virus in the next year.

This was our third lesson from the model, that *we could calculate the required amount of mosquito reduction*. In fact, we found that a 40–70% reduction in the New York City mosquito population would have prevented the 2000 outbreak. There is, of course, a catch with this approach: *the virus has to have already caused the initial mortalities that are used for subsequent control estimates*. However, as the virus continues to spread, bird mortality data are increasingly available. Perhaps the experience of cities and regions with the virus (and therefore with the data) can inform the control programs in areas where a virus outbreak has not yet occurred.

Like many research projects, this one raised more questions than it answered. How fast could West Nile virus spread across the continent and where would it go? How important was climate in determining the survival of birds, mosquitoes, and the virus? How would controlling the mosquitoes in an

urban area influence the surrounding rural area? How was our model similar or different compared to other models of similar diseases? Was there anything more general that could be learned about the epidemiology of vector-borne viruses? By working with some additional collaborators, we have begun to address these questions too.

In the meantime, I have an active interest in some very local mosquito control. Earlier this evening, I noticed that the spiders in the rocks behind my tent were making a killing, literally. Mosquitoes were landing in their webs so fast the spiders could hardly keep up. So now I'm starting to wonder how many spiders I would need to keep in my tent to control the mosquitoes. Or if not spiders, perhaps a bat would be more efficient? Clearly, I need another model. I'll have to consult with my collaborators.

Appendix

Bird equations:	
Susceptible	$\frac{dS_B}{dt} = -abI_M \frac{S_B}{N_B}$
Infectious	$\frac{dI_B}{dt} = abI_M \frac{S_B}{N_B} - \mu_V I_B - gI_B$
Removed	$\frac{dR_B}{dt} = (g + \mu_V) I_B$
Mosquito equations:	
Larval	$\frac{dL_M}{dt} = \beta_M(S_M + E_M + I_M) - mL_M - \mu_L L_M$
Susceptible	$\frac{dS_M}{dt} = -acS_M \frac{I_B}{N_B} + mL_M - \mu_A S_M$
Exposed	$\frac{dE_M}{dt} = acS_M \frac{I_B}{N_B} - kE_M - \mu_A E_M$
Infectious	$\frac{dI_M}{dt} = kE_M - \mu_A I_M$


Table 1: Ordinary differential equations (ODEs) representing the transmission of West Nile virus between birds and mosquitoes. Each equation keeps track of the change in a variable: the number of susceptible S_B , infected I_B , and removed R_B birds, and the number of larval L_M , susceptible S_M , exposed E_M , and infected I_M female mosquitoes per bird. Change in these variables is represented by parameters, which are probabilities and rates. (Note that all variables are indicated by capital letters, and all parameters by lower-case letters.) These equations match up with the boxes and arrows in Figure 1. To create your own numerical simulation of this model, you will need to choose the initial conditions, or values, for each variable. For a West Nile virus invasion, you might start with all variables equal to zero, except for I_B or I_M , which could be a very small positive number. Then you can set up the ODEs to change the values of the variables in each time step. What initial value of I_B or I_M do you need in order for the infection to increase? How do your results change if you allow birds to recover (i.e., set $g > 0$)?

Parameter	Mean (range)	Biological meaning
a	0.09 (0.03–0.16)	Mosquito daily per capita biting rate on birds
b	0.88 (0.80–1.00)	Probability per bite of West Nile virus being transmitted by an infectious mosquito biting a bird
c	0.16 (0.02–0.24)	Probability per bite of West Nile virus being transmitted by a mosquito biting an infectious bird
β_M	0.537 (0.036–42.5)	Mosquito daily per capita birth rate
m	0.068 (0.051–0.093)	Mosquito larva daily per capita maturation rate to susceptible adults
μ_A	0.029 (0.016–0.070)	Adult mosquito daily per capita mortality rate
μ_L	1.191 (0.213–16.9)	Larval mosquito daily per capita mortality rate
k	0.106 (0.087–0.125)	Mosquito daily per capita transition rate from exposed to infectious
μ_V	0.143 (0.125–0.200)	Bird daily per capita mortality rate from West Nile virus
g	0	Bird daily per capita recovery rate from West Nile virus

Table 2: Sample parameter values (and ranges) estimated from the literature for the West Nile virus model shown in Figure 1 and Table 1. The phrase *per capita* is a common biological term meaning per individual.

Further Information

- [1] This article is based on the paper: Wonham, M., T. de-Camino-Beck and M. Lewis, 2004: An epidemiological model for West Nile virus: Invasion analysis and control applications. *Proceedings of the Royal Society of London B*, **271**, 501–507.
- [2] Read more about this and other research at the University of Alberta's Centre for Mathematical Biology: <http://www.math.ualberta.ca/~mathbio/>.



Q: How does a mathematician induce good behaviour in his/her children?
A: "If I've told you n times, I've told you $n + 1$ times..."

Q: What keeps a square from moving?
A: Square roots.