

# **Appendix A from D. A. Kennedy et al., “Pathogen Growth in Insect Hosts: Inferring the Importance of Different Mechanisms Using Stochastic Models and Response-Time Data”**

**(Am. Nat., vol. 184, no. 3, p. 407)**

## **Supplementary Information**

### **Rearing Methods**

The first round of virus transmission in nature occurs when larvae emerge from surface-contaminated eggs (Doane 1969; Elkinton and Liebhold 1990). To ensure that no larvae became exposed to virus during hatch, we therefore surface-sterilized the egg masses by submerging them in 4% formaldehyde solution for 60 min, followed by rinsing in distilled water for 90 min (Dwyer and Elkinton 1995). Following standard protocols, eggs were left overnight to dry and were then placed in 6-oz plastic cups that were sealed with waxed paper lids. After treatment, eggs were stored at 25°C and 60% relative humidity to facilitate hatching. Two days later, larvae began to emerge from the eggs, and because we observed no virus-induced deaths, we concluded that the surface-sterilization was indeed successful.

Newly hatched larvae were transferred in groups of 30 to new 6-oz plastic cups, containing approximately 2 oz of a wheat-germ-based artificial diet (Bell et al. 1981). The larvae were reared at 25°C and 60% relative humidity. For our experiment we used only larvae in the fourth instar, because this instar is very important for disease dynamics in nature (Elkinton and Liebhold 1990). Resistance to virus changes dramatically over time even within instars (McNeil et al. 2010), and so to further reduce host heterogeneity, we synchronized development by harvesting third instars with head capsules that had begun to slip forward, a tell-tale signal that eclosion to the fourth instar is imminent. We then stored these larvae at 4°C for up to 72 h, a procedure that has been shown to slow development without altering susceptibility to infection (Hoover et al. 2002). This protocol allowed us to collect a large number of larvae that were developmentally synchronized at the period just before molting to the fourth instar. Once we had accumulated enough larvae for our experiment, we held the larvae for 48 h at room temperature to allow them to complete molting. The overall effect was that all larvae in our experiment were newly molted fourth instars.

### **Bioassay**

Following standard procedure for a diet plug bioassay (Hughes and Wood 1986; Dwyer et al. 1997; Li and Bonning 2007) we placed each larva in a 2-oz cup with a small plug of wheat germ diet (approximately 3 mm<sup>2</sup>). Each diet plug was contaminated with 3 mL of virus in distilled water, and the larvae were then allowed to feed overnight. We used five virus doses, created through serial dilutions of 2-, 2-, 4-, and 81-fold, respectively to create doses of 13,500, 6,750, 3,375, 844, and 10.4 viral occlusion bodies per larva, as well as a control solution that consisted of only distilled water. There were no virus-caused deaths in our control larvae. Larvae that consumed the entire diet plug were transferred to 2-oz cups containing virus-free diet, and larvae that failed to consume the entire plug were discarded (Dwyer et al. 2005). The remaining larvae were checked every 12 h until death or pupation and the times of death were recorded. Larvae that died because of the virus could usually be identified visually, but in any cases of uncertainty we examined smears under a light microscope to look for the polyhedral-shaped virus particles (Woods and Elkinton 1987). The quantity of virus released upon host death was not measured, because viral occlusion bodies pour out of larvae on death and mix with the diet. This quantity, however, has been well documented previously (Shapiro et al. 1986).

### **Fitting the Models to Data**

#### *Markov Chain Monte Carlo (MCMC)*

We begin by introducing the traditional Metropolis-Hastings MCMC algorithm. Importantly, the likelihood of a parameter set is proportional to the probability of the data given that parameter set:

$$L(\theta|D) \propto P(D|\theta), \tag{A1}$$

where  $L(\cdot)$  is the likelihood function,  $P(\cdot)$  is the density function,  $D$  is the observed data, and  $\theta$  is a vector of parameters.

Bayes' theorem says that:

$$P(\theta|D) = \frac{P(\theta)P(D|\theta)}{\int_{-\infty}^{\infty} P(\Theta)P(D|\Theta)d\Theta}, \quad (\text{A2})$$

where  $P(\theta | D)$  is the posterior density of the parameter set  $\theta$  given the data  $D$  and  $P(\theta)$  is the prior of  $\theta$ . By substitution with equation (A1), the above equation becomes

$$P(\theta|D) = \frac{P(\theta)L(\theta|D)}{\int_{-\infty}^{\infty} P(\Theta)L(\Theta|D)d\Theta}. \quad (\text{A3})$$

If the numerator and denominator on the right-hand side of equation (A3) can be calculated, it follows that the model posterior can be calculated. If only the numerator can be calculated, however, as is the case for our models and for many other nonlinear models, equation (A3) cannot be solved.

The premise of MCMC is that if a Markov chain were constructed with the appropriate transition probabilities, the chain would have a stationary state that is identical to the posterior distribution in equation (A3). Samples from this chain could then be used to estimate the posterior distribution of the model. To construct this Markov chain, it is necessary to define the acceptance probability:

$$\alpha(\theta_1, \theta_2) = \min\left(1, \frac{P(\theta_2|D)q(\theta_1)}{P(\theta_1|D)q(\theta_2)}\right). \quad (\text{A4})$$

Here,  $\alpha(\theta_1, \theta_2)$  is the probability of accepting a proposed jump from the current parameter set  $\theta_1$  to a new parameter set  $\theta_2$ , where  $\theta_2$  is drawn from an arbitrary proposal distribution  $q(\cdot)$ . Substituting from equation (A3), we can rewrite equation (A4) as

$$\alpha(\theta_1, \theta_2) = \min\left(1, \frac{P(\theta_2)L(\theta_2|D)q(\theta_1)}{P(\theta_1)L(\theta_1|D)q(\theta_2)}\right). \quad (\text{A5})$$

Since the proposal distribution  $q(\theta)$  can be defined arbitrarily and the prior distribution  $P(\theta)$  is defined by our prior knowledge of  $\theta$ , the likelihoods  $L(\theta_1 | D)$  and  $L(\theta_2 | D)$  are the only remaining terms necessary to calculate the acceptance probability  $\alpha(\theta_1, \theta_2)$ . We describe how these likelihoods were estimated in "Line-Search MCMC." Using the left-hand side of equation (A5) as an acceptance probability, we can then generate a Markov chain whose stationary state is equivalent to the posterior distribution of the model. The posterior distribution can thus be approximated by sampling from the stationary state of this Markov chain.

### Line-Search MCMC

Except for the Shortley birth-death model, our models are sufficiently complicated that there is no simple expression that gives the model's predicted distribution of response times. We must therefore approximate each model's predicted distribution through simulation. Indeed, part of the reason why Shortley (1965) considered only a linear birth-death model is because the solutions of that model can be calculated using a Bessel function, which can be approximated very accurately and rapidly without simulations.

Generation of even a single realization or "run" for the other models requires several seconds of computing time, while accurate estimation of a likelihood score for a parameter set requires hundreds to thousands of realizations. MCMC requires likelihood calculations for many thousands of parameter sets. Our model-fitting routines thus require a great deal of computing time, which in turn necessitated the use of a hybrid simulation algorithm. Specifically, at low population sizes, we used the Gillespie algorithm (Doob 1945; Gillespie 1977), which exactly simulates the underlying stochastic process (with the obvious exception of the deterministic model, for which we used a differential equation solver (Galassi et al. 2009)). For large population sizes, our stochastic simulations can be closely approximated by ordinary differential equations that track only the mean population size (fig. 2). Additionally, because the host immune system is usually exhausted early in an infection, by the time the virus population replicates to high levels, the rate of destruction of virus particles is close to zero, and so virus growth can be closely approximated by a logistic growth model (or exponential growth for the models that lack a virus carrying capacity). Accordingly, when the pathogen population size exceeded  $10^4$  virions, we used the analytic solution for the corresponding differential equation models to determine the remaining time to death. An alternative method to this hybrid simulation algorithm would be to use a tau leap algorithm (Cao and

Petzold 2006; Gillespie 2007). For this particular problem, our hybrid algorithm is computationally cheaper and still suitably accurate (“Hybrid Algorithm versus Tau-Leap Algorithm”).

To calculate the fit of the models to the data, we thus used model simulations to numerically evaluate the likelihood for a given parameter set. Each simulated trajectory represents one host exposed to the virus, ending either with complete clearance of the virus, in which case the host survives, or with the virus reaching the threshold population size that causes host death. In our experiment, we counted the number of insects that died every 12 h over 26 days. We therefore pooled our simulation results into discrete bins, with one bin representing hosts that survived, and 51 bins representing the 12-h time intervals between counts of the number dead. We then used these model realizations to parameterize a multinomial distribution according to the following equation:

$$p_i = \frac{2\nu_i + 1}{C + 2\sum_{j=1}^C \nu_j}, \quad (\text{A6})$$

where  $\nu_i$  is the number of model realizations that fell in bin  $i$ , and  $C$  is the total number of bins. Here, the sum of  $\nu_j$  over  $j$  is equal to the total number of model realizations. In practice  $\nu = 300$  was used in each likelihood evaluation for our parameter line searches, and  $\nu = 3,000$  was used in our MCMC chains. Our justification for choosing these values of  $\nu$  is given in “Likelihood Sensitivity.” We note that the values of  $p_i$  are adjusted slightly relative to the sample means to avoid likelihood scores of zero, but as the number of simulations increases, the effect of this adjustment becomes small. The likelihood of a particular parameter set can then be described by a multinomial distribution with  $p_i$  representing the probability of death in bin  $i$ . The likelihood of a parameter set  $\theta$  is then

$$L(\theta|d_1, d_2, \dots, d_C) = \frac{n!}{d_1!d_2! \dots d_C!} p_1^{d_1} p_2^{d_2} \dots p_C^{d_C}. \quad (\text{A7})$$

Here,  $\theta$  is a vector of parameters of the model,  $d_i$  is the number of hosts in the data that died during time bin  $i$ , and  $n$  is the total number of hosts used in the experiment. Note that because the probabilities  $p_1, p_2, \dots, p_C$  were generated from the model, they depend on the parameter set  $\theta$ . Also, because  $n$  is the total number of insects in the experiment, the sum of  $d_i$  over all  $i$  is  $n$ . This likelihood, the priors (defined in “Prior Construction”), and the proposal distributions (which we will describe shortly) can then be used according to equation (A5) to determine the acceptance probability of an MCMC chain. This MCMC chain is in turn used to estimate the posterior distribution of the model parameters.

The choice of the proposal distributions can substantially alter the rate at which MCMC chains converge on the posterior. Because simulations of our model require a great deal of computing time, we developed a modification of standard MCMC that we call “line-search MCMC,” which allowed us to design proposal distributions that roughly approximated the posterior distribution, leading to rapid convergence (Kennedy et al. 2014b). Here we briefly outline the method. First, we used a large number of parameter line searches, sometimes referred to as conditional optimizations, to find parameter sets with reasonably high posterior densities. In parameter line search, a single parameter is allowed to vary over a set range to find the conditional maximum posterior. The parameter is then fixed at the value that gives the maximum conditional posterior and the next parameter is likewise allowed to vary, until each parameter has been varied a set number of times. In practice, 4,000 parameter line searches were run for each model. Second, from these 4,000 searches, we selected the 50 parameter sets with the highest posterior density, and we carried out a principal components analysis on this subset of parameter sets. Third, we used the means and variances of these principal components as the means and variances of our proposal distributions. Using proposals based on principal components strongly reduced the correlations between successive parameter samples, which in turn led to rapid convergence on what appears to be the posterior distributions of our models.

Convergence of each model was determined using the Gelman-Rubin convergence diagnostic (Gelman and Rubin 1992) from the coda package (Plummer et al. 2009) in the R statistical language (R Development Core Team 2009). This diagnostic uses the output from multiple MCMC chains to compare the variance of the posterior distribution within chains to the variance of the posterior distribution between chains. Based on these variances, a summary statistic  $R$  is generated. This summary statistic approaches 1 as convergence is achieved. In practice, we ran five MCMC chains for each model, and we used a cutoff of  $R < 1.25$  to indicate convergence. We additionally examined trace plots to ensure that each chain was sampling from the same parameter space. These analyses yielded the posterior estimates shown in figure 5.

## Generalized Linear Model (GLM) Analyses

Fitting a GLM with link logit to the data showed that there was a statistically significant positive relationship between the log of virus dose and virus mortality (likelihood ratio test,  $\chi^2 = 791.733$ ,  $df = 1$ ,  $P < .001$ ), meaning that increasing the virus dose increased the likelihood of host death. A linear regression also showed that host response time significantly

decreased with the log of dose (likelihood ratio test,  $\chi^2 = 47.014$ ,  $df = 1$ ,  $P < .001$ ), meaning that increasing virus dose hastened the time to death of the hosts that died. Both of these observations are consistent with the predictions of our dynamical models (fig. 3).

## Hybrid Algorithm versus Tau-Leap Algorithm

As described in “Fitting the Models to Data,” we generated realizations of our model using a hybrid simulation algorithm. In this method, the Gillespie algorithm is used when virus population sizes are small, and the solution to the model’s ordinary differential equations is used to calculate the trajectory once the virus population reaches a predefined large size. An alternative method for simulating stochastic processes with large population sizes is the tau-leap algorithm (Cao and Petzold 2006; Gillespie 2007), in which state changes are determined through jumps forward in time using a Poisson approximation. Both methods reduce computing effort relative to the Gillespie algorithm, at the cost of approximating the underlying stochastic process. For our models, the hybrid algorithm provided a greater decrease in computing effort than did the tau-leap algorithm, which we implemented using the `adaptivetau` package (Johnson 2012) in the R programming language (R Development Core Team 2012). We thus used the hybrid algorithm in our model fitting.

We additionally examined whether our hybrid algorithm and the tau-leap algorithm gave similar results with regard to two important features of our data, time of death and percent mortality. For each algorithm, we simulated  $10^4$  realizations using our best-fit model (M2), assuming a virus dose of 844 virus occlusion bodies, and we sampled parameter sets with replacement from the posterior distribution shown in figure 5. Simulated times of death were grouped into bins using the binning method described in “Line-Search MCMC.” A  $\chi^2$  test then showed that there were no significant differences in mortality between the two algorithms ( $\chi^2 = 0.290$ ,  $df = 1$ ,  $P = .59$ ). The distribution of time to death did differ slightly between the two methods (fig. A1), in that our hybrid algorithm generated simulations with an average death time that was 19.41 h faster than the tau-leap algorithm. Given that we tallied mortality only every 12 h and that the mean kill time was about 280 h, this difference is probably trivial. We nevertheless explored the issue further in the following way. We note that until the population size reaches  $10^4$ , our hybrid algorithm uses the Gillespie algorithm, which is known to be exact. Any differences between the two algorithms in the time it takes to reach this  $10^4$  threshold must therefore be caused by approximation error in the tau-leap algorithm. We therefore used each method to simulate new trajectories as described above, except that we terminated each simulation when the virus population size reached  $10^4$ . This comparison revealed a mean difference of 3.40 h. Because this difference is entirely due to error in the tau-leap algorithm, we can conservatively say that 16.01 h or less of the mean difference between simulations with these two algorithms is caused by our hybrid algorithm.

## Prior Construction

Vague priors, such as those provided by normal distributions with large variances or by uniform distributions with large or even infinite ranges, are commonly used when there is little information regarding parameter values. Within-host virus growth models are almost never fit to data, and so one might expect that we would be forced to use vague priors. Because the gypsy moth baculovirus system is of economic and ecological interest, however, a great deal is known from experimental and observational studies, allowing us to construct informative priors for many parameters.

Single particles or “occlusion bodies” of the gypsy moth baculovirus include on the order of 100 virions (Ackermann and Smirnoff 1983; Shim et al. 2003). Previous studies, however, have suggested that the number of virus particles that actually cross the host gut is likely much smaller than the number of virions consumed (Zwart et al. 2009a). Our prior on the number of virus genomes that cross the gut per occlusion body consumed ( $c_1/[c_2 + 1]$ ) is therefore exponentially distributed with mean 100.

Previously, estimates of the replication rate  $\phi$  for baculoviruses of several Lepidopteran hosts other than the gypsy moth were directly measured in vivo (van Beek et al. 1990, 1988a). These estimates ranged from 0.242 to 0.394 replications per hour, and so our prior on  $\phi$  is gamma distributed with sufficient variance to include these values. We additionally note that van Beek and colleagues measured the net replication rate, which includes virus deaths due to the immune system, rather than the replication rate per se. For most of our models, the immune system is exhausted early in an infection, and so these two quantities rapidly become indistinguishable. The Shortley birth-death model (M3), however, has a constant death rate, and so for this model, we treat the quantity measured by van Beek and colleagues as the birth rate minus the death rate. The same gamma distribution described above for  $\phi$  in the other models was thus used for the prior of the net replication rate  $\phi - \beta m$  in the Shortley birth-death model.

Shapiro et al. (1986) used a light microscope to quantify the number of occlusion bodies released per infected fourth-instar gypsy moth larva. As we mentioned in the main text, we used the estimate from this study to set the median and standard deviation in a lognormal prior on the threshold number of virus particles that cause host death  $N$ .

Virus-induced larval death is marked by a rupturing of the host's cuticle and the release of massive quantities of occlusion bodies. With the exception of a few hard body parts such as the head capsule and setae, virus infection generally causes the majority of the host's tissue to be converted to virus. It is therefore reasonable to assume that the virus parasitizes most of the host cells suitable for replication and that the virus density at host death is close to its carrying capacity. As we discussed in the main text, it is convenient to express the virus carrying capacity  $\hat{K}$  as  $x_T K$ , where  $x_T$  is the number of virus particles released at death, and  $K$  is the amount by which the carrying capacity exceeds  $x_T$ . Because the carrying capacity  $\hat{K}$  is probably close to the threshold for death  $x_T$ , as a prior for  $K$ , we used an exponential distribution, because exponential distributions assume that the most likely values of  $K$  are small. We note that we first attempted to fit our models using an improper flat prior on  $K$ , but because of a lack of information about  $K$  in our data, this prior led to an improper posterior. We thus discarded the flat prior in favor of the exponentially distributed prior, which is better behaved and is biologically more reasonable.

There is very little information in the literature on the median number of immune cells present in a host,  $m$ , and the rate at which these immune cells destroy virus particles,  $\beta$ . For these latter parameters we therefore used vague half-normal priors, so that our estimates of these parameters depended almost entirely on our data. We note that although the initial number of immune cells  $y_0$  in any single host should take an integer value,  $y_0$  was drawn from a continuous distribution (eq. [6]). For the purposes of generating likelihoods,  $y_0$  was thus rounded to the nearest integer, except in the case of the no-demographic-stochasticity model (M7), which allows  $y_0$  to be continuous. It is, however, highly unlikely that this rounding step influenced our conclusions, because in our fitted models, realized values of  $y_0$  were large relative to the rounding error (fig. 5).

Formally, our priors are thus

$$p(\beta = z) = \begin{cases} \frac{2}{\sqrt{2\pi} \times 10^{14}} e^{-z^2/2 \times 10^{14}} & \text{if } z \geq 0 \\ 0 & \text{otherwise} \end{cases}, \quad (\text{A8})$$

$$p(m = z) = \begin{cases} \frac{2}{\sqrt{2\pi} \times 10^{14}} e^{-z^2/2 \times 10^{14}} & \text{if } z \geq 0 \\ 0 & \text{otherwise} \end{cases}, \quad (\text{A9})$$

$$\sigma_m \propto 1, \quad (\text{A10})$$

$$\frac{c_1}{c_2 + 1} \sim \text{exponential}(\lambda = 0.01), \quad (\text{A11})$$

$$\phi \sim \text{gamma}(k = 100, \theta = 0.00292), \quad (\text{A12})$$

$$N \sim \log - \mathcal{N}(\mu = \ln(2.049 \times 10^9), \sigma = 0.108), \quad (\text{A13})$$

$$\sigma_N \sim \log - \mathcal{N}(\mu = \ln(.25), \sigma = 0.1), \quad (\text{A14})$$

$$K - 1 \sim \text{exponential}(\lambda = 0.02), \quad (\text{A15})$$


where  $\propto 1$  denotes an improper flat prior over the set of nonnegative real numbers.

## Likelihood Sensitivity


Likelihood estimates evaluated from stochastic simulations can, in some cases, be inaccurate and imprecise. One solution to this problem is to use a large number of realizations in each likelihood evaluation. Because computing power is limited, a key question is therefore how many realizations are needed for a good estimate? For our models, we chose to use 300 realizations for each likelihood evaluation during our line searches and 3,000 realizations during our MCMC step. As we explained previously, the smaller number of realizations was used in the line-search step because many parameter sets provided very poor likelihoods that could be rejected easily. Moreover, the line searches were used only to develop proposal distributions for the subsequent MCMC steps, and so accurate estimates of the likelihood in this step were relatively unimportant.

Figure A2 shows how the likelihood estimate for our best parameter set changes with the number of realizations used. We note that the mean value of the likelihood changes as the number of realizations used increases, because of an adjustment term in equation (A6) that was introduced to avoid likelihood scores of zero. As shown in this figure, the effect of this adjustment disappears as the number of realizations increases. The standard deviation in our estimate of the likelihood likewise gets small as the number of realizations increases. Although there is some variation in the log

likelihood at 3,000 realizations (SD = 4.7), this variation decreases very slowly as we increase the number of realizations, such that SD = 2.5 at 30,000 realizations. Due to computational limits, we were thus forced to accept some error in our estimate of the likelihood. Our main conclusions, however, are based on model fits that differ in mean deviance by about an order of magnitude more than this error. We therefore have good reason to believe that a more precise estimate would not alter our conclusions.



**Figure A1:** The distributions of simulated response times generated using our hybrid algorithm (red squares) and using the tau-leap algorithm (circles). Each distribution is the result of  $10^4$  stochastic simulations. Each simulation is one realization of our best overall model, the linear virus growth model (M2), using parameter sets sampled with replacement from the posterior distribution of this model. Notice that although the two models predict similar distributions, the hybrid algorithm predicts on average slightly shorter times to death (19.41 h).



**Figure A2:** The relationship between the number of stochastic simulations used to evaluate the likelihood and the value of the log likelihood. Values plotted were evaluating using the maximum posterior parameter set for our best model (M2). Each point shows the mean of 300 log likelihood evaluations. Error bars show  $\pm 2$  SD.

## Literature Cited Only in Appendix A

- Ackermann, H., and W. Smirnov. 1983. A morphological investigation of 23 baculoviruses. *Journal of Invertebrate Pathology* 41:269–280.
- Bell, R., C. Owens, M. Shapiro, and J. Tardif. 1981. Mass rearing and virus production. *In* The gypsy moth: integrated pest management. USDA technical bulletin, Washington, DC.
- Cao, Y., D. Gillespie, and L. Petzold. 2006. Efficient step size selection for the tau-leaping simulation method. *Journal of Chemical Physics* 124:044109.
- Doane, C. 1969. Trans-ovum transmission of a nuclear-polyhedrosis virus in the gypsy moth and inducement of virus susceptibility. *Journal of Invertebrate Pathology* 14:199–210.
- Dwyer, G., and J. Elkinton. 1995. Host dispersal and the spatial spread of insect pathogens. *Ecology* 76:1262–1275.
- Dwyer, G., J. Firestone, and T. Stevens. 2005. Should models of disease dynamics in herbivorous insects include the effects of variability in host-plant foliage quality? *American Naturalist* 165:16–31.
- Galassi, M., J. Davies, J. Theiler, B. Gough, G. Jungman, P. Alken, M. Booth, and F. Rossi. 2009. GNU scientific library reference manual. 3rd ed., v1.12. Network Theory, Bristol, UK.
- Gillespie, D. 2007. Stochastic simulation of chemical kinetics. *Annual Review of Physical Chemistry* 58:35–55.
- Hoover, K., M. Grove, and S. Su. 2002. Systemic component to intrastadial developmental resistance in *Lymantria dispar* to its baculovirus. *Biological Control* 25:92–98.
- Johnson, P. 2012. Adaptivetau: tau-leaping stochastic simulation. R package version 1.0. <http://cran.r-project.org/package=adaptivetau>.
- R Development Core Team. 2012. R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna.
- Shim, H., J. Roh, J. Choi, M. Li, S. Woo, H. Oh, K. Boo, and Y. Je. 2003. Isolation and characterization of a *Lymantria dispar* multinucleocapsid nucleopolyhedrovirus isolate in Korea. *Journal of Microbiology* 41:306–311.
- van Beek, N., P. Flore, H. Wood, and P. Hughes. 1990. Rate of increase of *Autographa californica* nuclear polyhedrosis virus in *Trichoplusia ni* larvae determined by DNA-DNA hybridization. *Journal of Invertebrate Pathology* 55:85–92.
- Woods, S., and J. Elkinton. 1987. Bimodal patterns of mortality from nuclear polyhedrosis-virus in gypsy-moth (*Lymantria-dispar*) populations. *Journal of Invertebrate Pathology* 50:151–157.