Supplementary materials for this article are available at https://doi.org/10.1007/s13253-023-00581-y.

Predicting the Temperature-Driven Development of Stage-Structured Insect Populations with a Bayesian Hierarchical Model

Kala STUDENS, Benjamin M. BOLKER, and Jean-Noël CANDAU

The management of forest pests relies on an accurate understanding of the species' phenology. Thermal performance curves (TPCs) have traditionally been used to model insect phenology. Many such models have been proposed and fitted to data from both wild and laboratory-reared populations. Using Hamiltonian Monte Carlo for estimation, we implement and fit an individual-level, Bayesian hierarchical model of insect development to the observed larval stage durations of a population reared in a laboratory at constant temperatures. This hierarchical model handles interval censoring and temperature transfers between two constant temperatures during rearing. It also incorporates individual variation, quadratic variation in development rates across insects' larval stages, and "flexibility" parameters that allow for deviations from a parametric TPC. Using a Bayesian method ensures a proper propagation of parameter uncertainty into predictions and provides insights into the model at hand. The model is applied to a population of eastern spruce budworm (*Choristoneura fumiferana*) reared at 7 constant temperatures. Resulting posterior distributions can be incorporated into a workflow that provides prediction intervals for the timing of life stages under different temperature regimes. We provide a basic example for the spruce budworm using a year of hourly temperature data from Timmins, Ontario, Canada.

Supplementary materials accompanying this paper appear on-line.

1. INTRODUCTION

Many efforts have been made to improve predictions of insect development and phenology in the past century (Uvarov 1931). These studies were historically motivated by the management of insect pests (Pruess 1983; Crimmins et al. 2020). More recently, climate change and the increasing threat from invasive species have renewed interest in incorporating quantitative models of insect development and phenology in process-based models of

K. Studens (✉)· J.-N. Candau, Natural Resources Canada, Canadian Forest Service, Great Lakes Forestry Centre, Sault Ste. Marie, ON, Canada. (E-mail: *kala.studens@nrcan-rncan.gc.ca*).
B. M. Bolker, Departments of Mathematics & Statistics and Biology, McMaster University, Hamilton, ON, Canada.

Journal of Agricultural, Biological, and Environmental Statistics
https://doi.org/10.1007/s13253-023-00581-y

Published online: 16 November 2023

geographic distribution (Régnière et al. 2012), evolutionary ecology (Bewick et al. 2016), and biological invasion (Porter et al. 1991).

Insects accomplish their life cycle by developing through discrete morphological stages. Development rate depends on several climatic variables, of which temperature is considered the most important (Rebaudo and Rabhi 2018). Most models of insect development rely on two principles: a *thermal performance curve* (TPC), which formalizes the relationship between development rate and temperature (Chuine and Régnière 2017), and a *rate summation*, which states that maturation to the next stage occurs when development (based on the TPC and the distribution of environmental temperature) reaches a threshold.

Bayesian methods have been applied to fit TPCs for various organisms. These include thermal development curves for unicellular organisms (Corkrey et al. 2012), temperature-dependent growth of fish (Childress and Letcher 2017) and fungi (Gajewski et al. 2021), and temperature-trait relationships of disease-vector insects (Johnson et al. 2015; Mordecai et al. 2019; Shocket et al. 2020; Villena et al. 2022). To our knowledge, Bayesian methods have never been used to model TPCs of forest insects. In this domain, TPC models fits typically provide point estimates with standard errors or confidence intervals (CIs: Rebaudo and Rabhi 2018; Quinn 2017); estimates of parameter correlation are rarely given. Most models of insect development are fitted using either nonlinear least squares or maximum likelihood estimation (MLE: Damos and Savopoulou-Soultani 2012). While one can use these methods to estimate correlations among parameters and hence derive CIs for prediction (Bolker 2008), Bayesian methods provide samples of the full multivariate posterior distribution and thus enable CIs for any function of model parameters. Bayesian methodology also allows the use of the vast amount of prior knowledge gathered on insect phenology [e.g. Hoffmann et al. (2013), Rebaudo and Rabhi (2018)] to specify informative prior distributions (McCarthy and Masters 2005). McManis et al. (2018) impose Bayesian prior distributions on the parameters of their mountain pine beetle development model, but the parameter estimation is done by minimizing the negative log posterior rather than sampling a full posterior distribution. This approach leverages prior knowledge, but does not provide the benefits of convenient error propagation.

In this study we developed a Bayesian hierarchical model of individual-based, temperature-driven development for spruce budworm (*Choristoneura fumiferana*). The resulting posterior distribution was used to simulate one year of larval development in the wild, to demonstrate how we could obtain prediction intervals for development achieved by a given date, or of dates by which a critical developmental threshold occurs.

2. METHODS

2.1. STUDY SYSTEM

The eastern spruce budworm is native to the boreal forests of North America. During periodic outbreaks, its populations have caused extensive tree mortality, especially of white spruce and balsam fir trees (Blais 1983). The species is univoltine, and its six larval stages (*instars*) are delineated by the shedding of head capsules. Upon emergence from eggs in the summer, the first instar larvae disperse and form hibernacula (cocoon-like shelters) in

which they moult to second instar and enter diapause, a resting stage in which they spend the duration of the winter. Second instar larvae emerge from these hibernacula in the early spring and feed on old needles until budburst. In the third to sixth larval instars, larvae feed on new foliage. The synchrony of these larval stages with new foliage is crucial for survival (Lawrence et al. 1997). Understanding the timing of each phenomenon is necessary in order to predict the impact of climate change on the species' behaviour. The pupal stage is typically reached in early summer, and moths emerge several days later. Understanding the timing of adult emergence is important for modelling the species' landscape-scale dispersal (Sturtevant et al. 2013).

Early models of spruce budworm development neglected individual variability (Bean 1961; Cameron et al. 1968; Dennis et al. 1986; Lysyk 1989; Hudes and Shoemaker 1988). Régnière (1984; 1987) was the first to account for variability in a two-step procedure that separately estimated the effect of temperature on median development rate and the distribution of individual rates around the median, modelling the rate using a Type I generalized logistic distribution (Balakrishnan and Leung 1988). Stedinger et al. (1985) modelled the distribution of larvae in each larval stage as Dirichlet-multinomial distribution, accounting for environmental conditions at three spatial scales (regional, site and individual). More recently, Régnière et al. (2012) proposed a general framework for quantifying insect development that models individual variability as a log-normal distribution.

2.2. DATA COLLECTION

The data come from a spruce budworm rearing experiment described in Wardlaw et al. (2022). Samples of diapausing larvae were collected as described in Candau et al. (2019) from Timmins, Ontario, in accordance with the methods in Perrault et al. (2021). Upon emergence from diapause, individuals were collected and placed in separate containers containing artificial diet created to mimic the nutrition that the insects would consume in the wild. The development of each insect was observed daily; a moult was reported once a larva had shed its head capsule. The colony was divided into seven sub-populations, each of which was reared at a different constant temperature, at evenly spaced temperatures from 5 to 35 °C. Due to the potential for high mortality rates at temperatures outside a "sustainable" temperature range of 15–25 °C, an extra rearing step was taken for populations at the extreme temperatures. For each stage, the times to moult at the extreme temperatures were estimated using the development rate model presented in (Régnière et al. 2012). The insects were held at the extreme temperatures for approximately half the predicted moulting time and then moved to 20 °C for the remainder of each stage. Each sub-population contained 250 individuals, resulting in a total of 1750 individuals. Since the main objective of the project was to observe development rates in the larval instar stages, survival was not considered in the modelling process. Thus, the data used to fit the model consisted only of individuals who survived the full rearing process from the first larval instar to pupation. These data included 76 individuals reared at 5 °C, 151 individuals at 10 °C, 18 at 15 °C, 120 at 20 °C, 92 at 25 °C, 197 at 30 °C and 109 individuals at 35 °C. Only the first new generation from the wild population was used to fit the data, to eliminate any generational effect of laboratory rearing. The low survival rate in the 15 °C treatment is likely due to the fact that it is on the

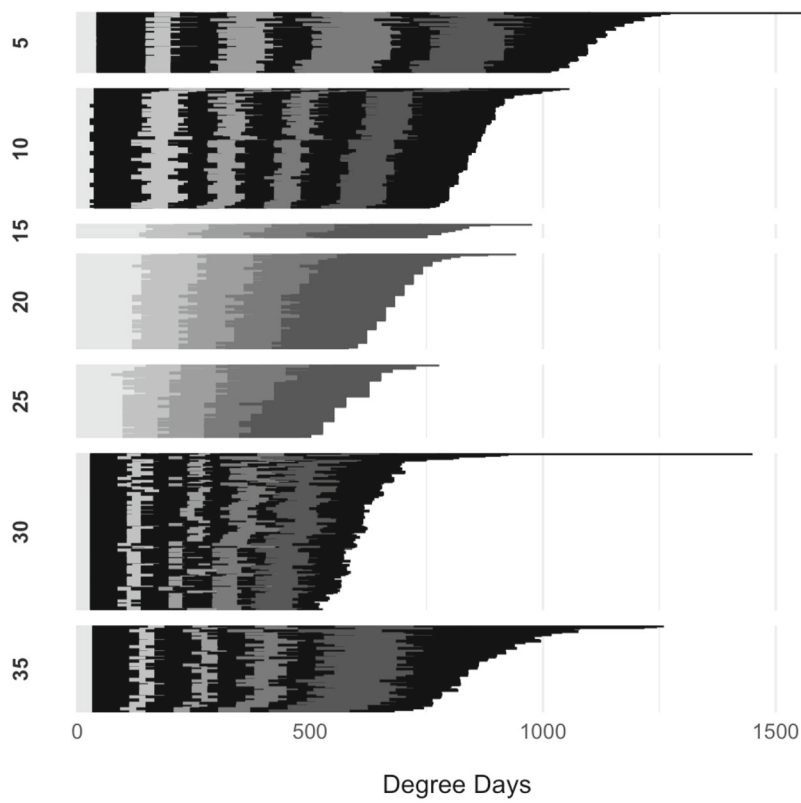

Figure 1. Visualization of the data used to fit the model, faceted by treatment. Grey bars indicate time spent at the treatment temperature with the shade darkening as stages progress, while black bars represent time spent at our alternative sustainable temperature, 20 °C. Plot facet heights are scaled by the number of observations in each treatment .

edge of the sustainable growth regime. The insects at this temperature may have benefited from the transfer treatment performed for the larvae exposed to more extreme temperatures (Fig. 1).

2.3. DEVELOPMENT RATE MODEL

We chose the Schoolfield et al. (1981) model with the reference temperature generalization presented in Ikemoto (2005) for our Bayesian implementation of the likelihood framework described in Régnière et al. (2012). The development rate equation in this model, representing the rate of the enzyme-catalysed reaction governing growth, is

$$r(T_K) = \frac{\rho \frac{T_K}{T_A} \exp\left[\frac{H_A}{R}\left(\frac{1}{T_A} - \frac{1}{T_K}\right)\right]}{1 + \exp\left[\frac{H_L}{R}\left(\frac{1}{T_L} - \frac{1}{T_K}\right)\right] + \exp\left[\frac{H_H}{R}\left(\frac{1}{T_H} - \frac{1}{T_K}\right)\right]} \quad . \tag{1}$$

here $R = 1.987\,\text{kcal} \cdot \text{K}^{-1} \cdot \text{mol}^{-1}$ is the universal gas constant and T_K is the independent temperature variable of the TPC in degrees Kelvin, while the remaining values are the

curve parameters. All temperatures in this model are represented in degrees Kelvin. This model formulation is based on the assumption that the enzymes controlling growth rate can be in one of three states: active, inactive due to low temperature, or inactive due to high temperature. These states are assumed to be reversible. The reciprocal of the denominator represents the fraction of enzymes in the active state. T_A is the temperature at which the fraction of active enzymes is maximized, while the value of ρ in the numerator represents the development rate when $T_K = T_A$. In the parameterization from Ikemoto (2005), the value of the reference temperature T_A is calculated from the other parameters as follows:

$$T_A = \frac{H_L - H_H}{R \cdot \log\left(\frac{-H_L}{H_H}\right) + \left(\frac{H_L}{T_L}\right) - \left(\frac{H_H}{T_H}\right)} \tag{2}$$

The parameters T_L and T_H represent the temperatures at which half of the rate-controlling enzymes have been de-activated due to low or high temperatures, respectively. The H_L and H_H parameters represent the changes in enthalpy resulting from low- and high-temperature enzyme de-activation, and H_A is the activation enthalpy of the growth reaction.

2.4. LIKELIHOOD STRUCTURE

In the model of Régnière et al. (2012), an individual's physiological age a within a developmental stage is defined as the proportion of the stage that it has completed. The age of an individual i in the development stage s is represented as:

$$a_{is}(t, \boldsymbol{\theta}_s) = \delta_{is} \sum_t r(T_t, \boldsymbol{\theta}_s)\Delta_t, \tag{3}$$

where $r(T, \boldsymbol{\theta}_s)$ is the population median development rate as a function of temperature T and the vector of development rate parameters $\boldsymbol{\theta}_s$ for stage s, and δ_{is} is a multiplier representing the development rate of individual i, relative to the population median for stage s. Here, δ_{is} is independent of temperature. The multiplier Δ_t represents the time spent at temperature T_t. The value a_i represents the proportion of the development stage s that has been completed by individual i, and thus, it ranges between 0 and 1.

An age of 1.0 signals the end of a developmental stage. Therefore, if t_m is the amount of time taken by individual i to complete a stage s, then $a_{is}(t_m, \boldsymbol{\theta}) = 1$. Since observations were taken daily, we do not directly observe t_m and instead the data are interval censored. That is, we know only the one-day interval within which an individual completed each developmental stage (i.e. $t_m \in [t_1, t_2]$). Thus,

$$a_{is}(t_1, \boldsymbol{\theta}) \; < \; a_{is}(t_m, \boldsymbol{\theta}) = 1 \; < \; a_{is}(t_2, \boldsymbol{\theta}). \tag{4}$$

Substituting for Eq. 3 and rearranging gives

$$\frac{1}{\sum_{t=1}^{t_2} r(T_t, \boldsymbol{\theta}_s)\Delta_t} < \delta_{is} < \frac{1}{\sum_{t=1}^{t_1} r(T_t, \boldsymbol{\theta}_s)\Delta_t}. \tag{5}$$

If we treat δ_{is} as a random variable with cumulative distribution F_δ, then the likelihood of our observation of individual i for stage s given parameters $\boldsymbol{\theta}_s$ is

$$\mathcal{L}(\boldsymbol{\theta}_s) = F_\delta\left(\frac{1}{\sum_{t=1}^{t_1} r(T_t, \boldsymbol{\theta}_s)\Delta_t}\right) - F_\delta\left(\frac{1}{\sum_{t=1}^{t_2} r(T_t, \boldsymbol{\theta}_s)\Delta_t}\right). \tag{6}$$

Régnière et al. (2012) assign a log-normal distribution to δ_{is} to ensure that $\delta > 0$ and that the distribution of individual variation is positively skewed. Their parameterization gives $E(\delta) = 1$, thus centring the population mean at the development rate curve. In our parameterization, we centre the distribution at the population median instead of the mean, so that a population's development times and development rates have the same distribution across individuals. Therefore, we have

$$\log(\delta_{is}) = \epsilon_{is} \overset{\text{iid}}{\sim} \text{Normal}(0, \sigma_{\epsilon_s}^2) \tag{7}$$

for each individual $i = 1, ..., N$, where σ_{ϵ_s} is unknown.

For each treatment j and stage s, Régnière et al. (2012) included a multiplicative random effect υ_{js}, which allows the mean development times at the treatment temperatures to deviate from those predicted by the parametric model. This term adds flexibility to the model, relaxing the strict parametric form of the development rate curve. These additional flexibility parameters are included to decrease the systematic error caused by misspecification of the development rate curve. Using this method allows us to obtain the most accurate predictions possible while still learning about the parameters from the model form that we are interested in.

Régnière et al. (2012) specified $\upsilon_{j,s} \sim \text{Normal}(1.0, \sigma_{\upsilon_s}^2)$ (where $\sigma_{\upsilon_s}^2$ is a fitted parameter). We set $\upsilon_{j,s} \sim \text{log-Normal}(0, \sigma_{\upsilon_s}^2)$ to ensure that the $\upsilon_{j,s}$ parameters remain nonnegative. The resulting development rates are:

$$r^*(T_j, \boldsymbol{\theta}_s) = \frac{r(T_j, \boldsymbol{\theta}_s)}{\upsilon_{js}} \tag{8}$$

Since we only have observations at seven treatment temperatures, we can only apply our flexibility parameters at those discrete points on the TPC. To expand this relaxation of the TPC's parametric form for prediction, we use the `interpSpline` function from the `splines` R package, to connect these, and the points (r, T) = (0, 0) and (40, 0) (i.e. we know development should approach zero towards these temperatures) so that we can estimate development at temperatures other than those observed in the experiment (Fig. 2).

The full log likelihood for stage s can be written as:

$$\ell(\boldsymbol{\theta}_s, \boldsymbol{\upsilon}) = \sum_j \left\{ \log\left[f_{\upsilon_s}(\upsilon_{js})\right] + \sum_i \log\left[F_{\delta_s}\left(\frac{1}{\sum_{t=1}^{t_{1_i}} r^*(T_{j_t}, \boldsymbol{\theta}_s)\Delta_t}\right) - F_{\delta_s}\left(\frac{1}{\sum_{t=1}^{t_{2_i}} r^*(T_{j_t}, \boldsymbol{\theta}_s)\Delta_t}\right)\right]\right\} \tag{9}$$

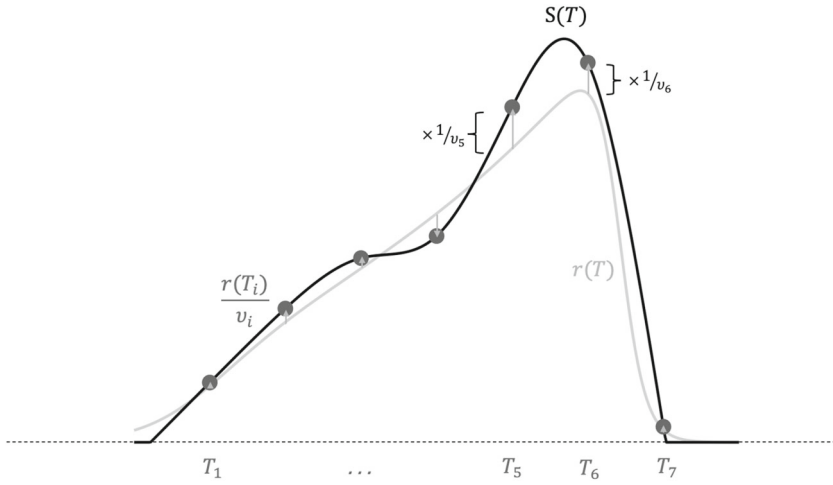

Figure 2. Example of the effect of flexibility parameters, υ, with modified development rates connected with an interpolation spline (R's `interpSpline` function).

where i denotes individual, j denotes temperature treatments, and f_{υ_s} is the probability density function of a log-normal distribution with location parameter 0 and scale parameter σ_{υ_s}.

The observations for the L2 stage were interval censored, while the observations for the remaining stages were double censored (Calle 2002). That is, the start time for the L2 stage is known, but the exact end time is interval-censored, while for the remaining stages both the start and end times are interval-censored. Since observations were made daily, the interval for the true moult time t for an individual in the L2 stage is $t \in [t_{\text{obs}} - 1, t_{\text{obs}}]$, while the interval for an individual in any other development stage is $t \in [t_{\text{obs}} - 1, t_{\text{obs}} + 1]$. Here, t_{obs} is the observed moult time. For the L2 stage, it is the number of days until a moult is observed, while for the remaining stages it is the number of days between the observations of successive moults (Fig. 3).

2.5. BAYESIAN IMPLEMENTATION

Effective methods for sampling from the posterior distribution of Bayesian models have been an active area in computational statistics for decades. Hamiltonian (or hybrid) Monte Carlo (HMC) proposed that nearly 30 years ago (Neal 1993) has recently achieved much greater popularity with the availability of convenient and powerful implementations. While Stan (Stan Development Team 2020) is probably the most popular such tool, our model was written in Template Model Builder (Kristensen et al. 2015), which is primarily designed for maximum likelihood estimation. The R package `tmbstan` (Monnahan and Kristensen 2018) was used to apply the No-U-Turn Sampler (NUTS), a particular step-size rule for HMC, to the model.

The chains were initialized by sampling parameter vectors from the prior distribution. We assumed Gamma priors for ρ_s (for $s = 1, ..., 5$), H_A, H_H, $-H_L$, and tdiff $= T_H - T_L$ to ensure nonnegativity, while T_L was assigned a log-normal prior. The scale parameters

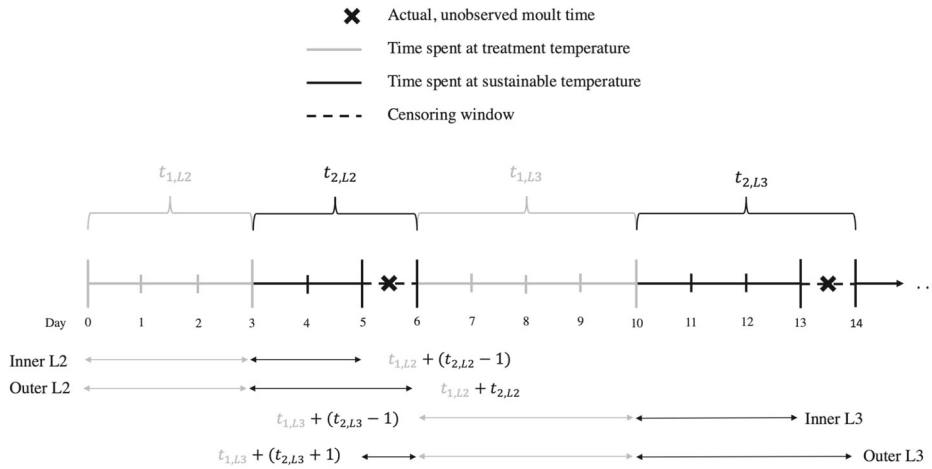

Figure 3. Example of interval censoring for L2 stage, and double interval censoring for L3 stage in a transfer treatment.

σ_{ϵ_s} and σ_{υ_s} were also assigned log-normal priors. We know that the duration of the spruce budworm's larval stages decreases from the second larval stage to the fourth and then increases until pupation (Régnière et al. 2012). Thus, while maintaining a constant scale parameter across stages, the shape parameters were set such that the means of the prior distributions for ρ reflected this known structure. The parameter ρ is the only TPC parameter that is permitted to vary across development stages, since we know that the stage durations vary. The other parameters are held constant under the assumption that the enzyme activation/deactivation process is the same across developmental stages. Our priors were determined both by interpreting the biological parameters and performing prior predictive checks. We plotted development rate curves with wide but realistic ranges on each parameter and tightened the prior assumptions when the resulting development rate curves displayed unrealistic characteristics (such as non-negligible development rates at temperatures above or below the range known to support development). To test the convergence and reliability of MCMC sampling, we used several diagnostic tools: the $\hat{R}$ convergence diagnostic and bulk/tail effective sample size (ESS) values (Vehtari et al. 2021), checking for divergent transitions (Stan Development Team 2020), simulation-based calibration (SBC: Talts et al. 2018; Cook et al. 2006) and posterior predictive checking (PPC). SBC entails several instances of sampling from the prior, generating a dataset using the model's data generating process, fitting the model to each of these datasets and assessing the uniformity of the resulting ranks of the prior draws within the posterior distributions. In theory, these ranks should be uniform. Visual tests are performed on the histograms of the ranks to check for uniformity. The full details and results of the SBC can be found in the Supplementary Materials.

Algorithm 1 Model Specification

$\rho_s \sim \text{Gamma}(k = -1.05s^2 + 4.22s + 4.08, \theta = 0.0045)$, for $s = 1, ..., 5$
$(-H_L, H_A, H_H) \sim \text{Gamma}(k = (3.6, 5.4, 7.6), \theta = (2.3, 3.1, 0.1))$
$T_L \sim \text{logNormal}(\mu = 5.64, \sigma = 0.0067)$
$T_H - T_L \sim \text{Gamma}(k = 112, \theta = 0.23)$
$(\sigma_\epsilon, \sigma_\upsilon) \sim \text{logNormal}(\mu = (-1.5, -2.5), \sigma = (0.1, 0.05))$
$\upsilon^*_{j,s} \sim \text{Normal}(\mu = 0, \sigma = 1)$, for $j = 1, ..., 7$ and $s = 1, ..., 5$
$LL \leftarrow -\log\left[\pi(\tilde{\boldsymbol{\theta}})\right]$
for i in 1 : nrow(data) **do** ▷ Here, $r(T, \boldsymbol{\theta})$ is the Schoolfield development rate model.
  $T_A \leftarrow (H_L - H_H)\left[R\log\left(-\frac{H_L}{H_H}\right) + \frac{H_L}{T_L} + \frac{H_H}{T_H}\right]^{-1}$
  $p_k \leftarrow r(T = t_k(i), \rho_{s(i)}, H_A, T_L, H_L, T_H, H_H, T_A)$ ▷ for $k = 1, 2$
  $\upsilon_k \leftarrow \exp\left[\upsilon^*_{t_k[j(i)],s}\sigma_\upsilon\right]$
  $p^*_k \leftarrow \frac{p_k}{\upsilon_k}$
  $\epsilon_{\text{lower}} \leftarrow \log(t_{1d}(i)p^*_1 + t_{2d}(i)p^*_2)\sigma^{-1}_{\epsilon,s}$
  $\epsilon_{\text{upper}} \leftarrow \log(t_1(i)p^*_1 + t_2(i)p^*_2)\sigma^{-1}_{\epsilon,s}$
  $ll(i) \leftarrow \Phi(\epsilon_{upper}) - \Phi(\epsilon_{lower})$
  $LL \leftarrow LL - n(i)ll(i)$
end for

2.6. POSTERIOR SIMULATIONS

To demonstrate how we could use our model fits to obtain downstream credible intervals for values that are of use for pest-management decision making, we simulated insect development based on our posterior samples and on a year of hourly weather data from Timmins, Ontario, where the wild colony was initially sampled. These data were obtained from the `weathercan` package in R (LaZerte and Albers 2018). Since the scope of the model did not include overwintering larvae, each individual was assumed to begin the L2 stage at the same time, on April 1, 2020. This procedure is outlined in Algorithm 1.

The data and the code for the model fitting, SBC and weather simulations can be found at https://github.com/kdis19/bayessbw.

3. RESULTS

Some divergent transitions did occur in earlier stages of the model development, but we eliminated them by (1) reducing the sampler's step size and (2) using non-centred parameterizations for distributions in the model definition in order to make the posterior surface more amenable to HMC sampling (Gelman et al. 2020; Stan Development Team 2020). Non-centred parameters are sampled from a standardized distribution and then shifted and scaled appropriately, rather than being sampled directly (Stan Development Team 2020). While the $\hat{R}$ diagnostic and ESS values were acceptable and the HMC sampler reported no divergent transitions, the results of the SBC indicated that the sampler was not properly recovering the scale parameters for the distribution of individual variation, σ_{ϵ_s}. In an attempt to correct these estimates for use in prediction, we fitted parameters to power-transform the

Algorithm 2 Posterior Simulations

for p in $1 : P$ **do** ▷ $P = 100$ populations
  $\boldsymbol{\theta}_p \sim f(\boldsymbol{\theta}|y)$ ▷ Sample from posterior
  $r_s^*(T_t, \boldsymbol{\theta}_p) \leftarrow r_s(T_t, \boldsymbol{\theta}_p)/\upsilon_{s,t}$
  Create a cubic spline connecting $r_s^*(T_t, \boldsymbol{\theta}_p)$ evaluated at each treatment temperature T_t using R's `interpSpline` function from the `splines` package:

$$S_s(T) = \begin{cases} P_{s,0}(T) & \text{if } T < 5 \\ P_{s,1}(T) & \text{if } 5 \leq T < 10 \\ \vdots & \\ P_{s,7}(T) & \text{if } 30 \leq T < 35 \\ P_{s,8}(T) & \text{if } T \geq 35 \end{cases}$$

  where $P(T_t) = r^*(T_t, \boldsymbol{\theta}_p)$
  for i in $1 : N$ **do** ▷ $N = 100$ individuals
    $\epsilon_{s,i} \sim \text{Normal}(0, \sigma^2_{\epsilon,p})$ ▷ $\sigma_{\epsilon,p}$ is an element of $\boldsymbol{\theta}_p$
    $\delta_{s,i} \leftarrow e^{\epsilon_{s,i}}$ ▷ Both for $s = 1, ...5$

    Calculate age vector using hourly weather data:
    $a_{s,i}(j) \leftarrow 0$ for $j = 1, ..., H$ and $s = 1, ..., 5$
    $s \leftarrow 1$
    for h in $1 : H$ **do** ▷ H is the length of the hourly weather vector
      $a_{new} \leftarrow a_{s,i}(h-1) + \delta_{s,i} S_s(T_h) \Delta h$
      if $a_{new} \geq 1$ **then**
        $a_{s,i}(j) \leftarrow 1$ for $j \geq h$
        if $s < 5$ **then**
          $s \leftarrow s + 1$
          $a_{s,i}(h) \leftarrow \delta_{s,i} S_s(T_h) \cdot \frac{a_{new}-1}{\delta_{s-1,i} S_{s-1}(T_h)}$
        else
          Break loop
        end if
      else
        $a_{s,i}(h) \leftarrow a_{new}$
      end if
    end for
    The age a of an individual i at the h^{th} time step is defined as $a_{i,h} \leftarrow \sum_s a_{i,s}(h)$
  end for
end for
Calculate temperature-wise prediction intervals for insect age using age vectors of all simulated individuals.

posterior draws from the SBC such that they matched with the prior. This transformation improved the coverage of 90% credible intervals for the σ_ϵ parameters.

We applied this fitted transformation to the posterior distribution we obtained by fitting the model to our data in order to improve the accuracy of our predictions. While this matched the prior with the full set of posterior samples, the SBC histograms still showed signs of heavy autocorrelation, as indicated by large spikes at the extremes of the histogram (Talts et al. 2018). The histograms for the remaining parameters looked mostly uniform with

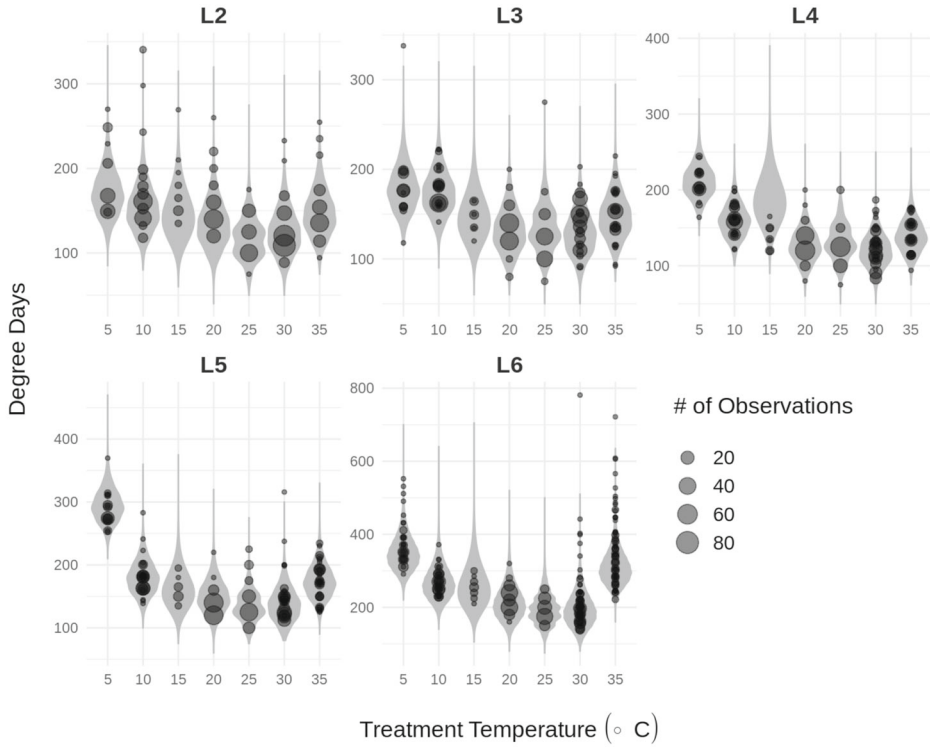

Figure 4. Violins represent densities of posterior prediction intervals for the number of degree days at each treatment temperature in each developmental stage. The points are the observed degree days at each corresponding stage and temperature.

some indications of skewness, but the boxplots comparing prior and posterior samples were very similar and the coverage of the 90% credible intervals ranged from 0.84 to 0.91 with mean 0.87. However, the SBC histograms for the development rate estimates at each rearing temperature and stage also displayed this tendency, while the boxplots indicated that the overall posterior and prior distributions were very similar.

Figure 4 compares posterior prediction intervals for the "observed" degree days to moult to the actual degree days observed in the data. The posterior datasets were generated in a similar fashion to Algorithm 1, where the weather vector inputs imitate the experimental conditions and the observations are interval censored as though they were observed under our experimental protocol. Degree days were used instead of days in this plot so that we could directly compare data from transfer treatments to that from non-transfer treatments. To generate the prediction intervals, we drew 1000 posterior samples containing 250 individuals per treatment. The data mostly fall well within the posterior prediction intervals, but there are a few outlier points that land outside (L6, 30 and 35 °C). These outlying points could be due to the discrepancy we observed with the σ_ϵ parameters in the SBC.

The bottom panel of Fig. 5 shows quantiles of daily development for the full simulated cohort. The vertical extent of the ribbons represents the credible intervals for the cohort's developmental age over time, while the ribbons' horizontal extents show the credible inter-

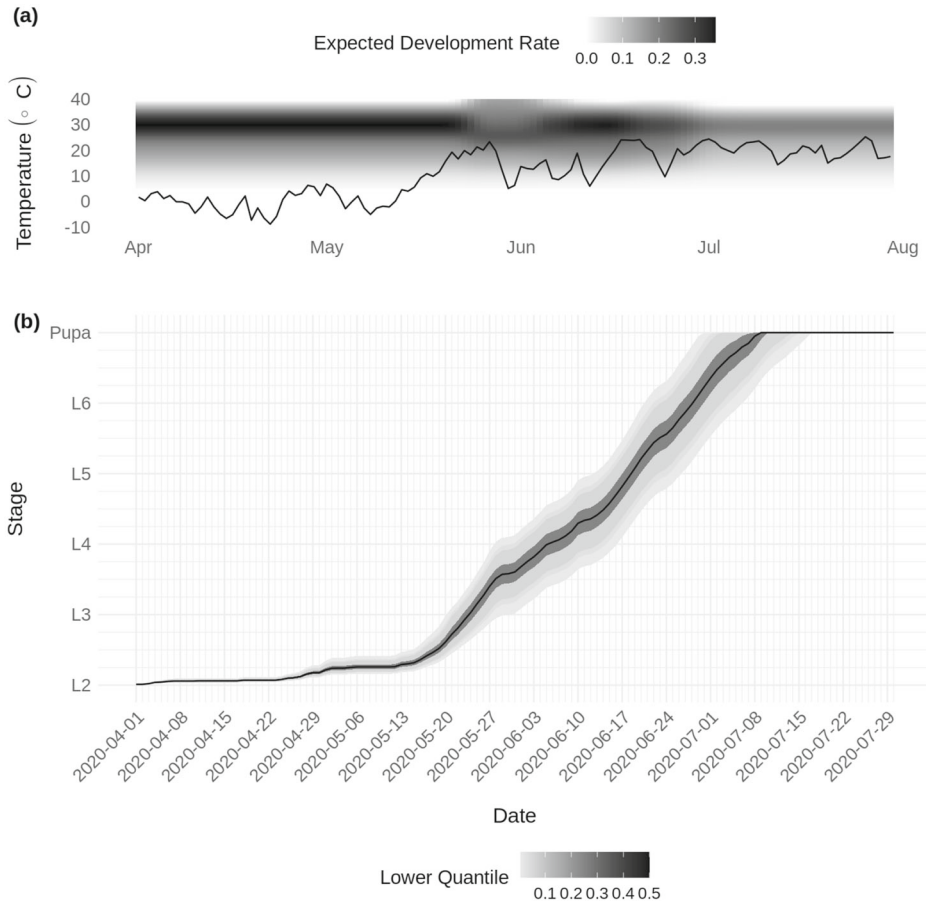

Figure 5. Top panel show mean daily temperature in Timmins, ON, for the year 2020. Background shading indicates expected development rate as a function of temperature and median proportions of the population in each development stage. Bottom panel shows symmetric credible intervals of development for the simulated population. The colours of the ribbons represent the lower percentile of the credible intervals .

vals of dates at which developmental milestones are reached. Though it uses a crude estimate of the cohort's starting date for the L2 stage, the development timeline shown in the figure matches what we would expect from an Ontario budworm population in a typical year (Régnière et al. 2012). The black line in the upper panel of the figure shows mean daily temperature, while the background shading represents the expected development rate given temperature and proportion of individuals in each stage at a given date. This expected development rate is calculated using the median estimated development rate curve for each stage, weighted by the median estimated proportion of individuals in each stage on a given date. When the mean daily temperature line traverses darkly shaded areas (e.g. early June), budworm larvae are expected to develop rapidly.

4. DISCUSSION

4.1. MODELLING ASSUMPTIONS

Several assumptions have to be made when accepting the results of the weather simulations. While our colony was collected from the wild, it was reared under laboratory conditions on artificial diet, at constant temperatures. While this is convenient for model fitting, we acknowledge that the results may not apply directly to populations of spruce budworm living in the wild [but see von Schmalensee et al. (2021)]. This model should be validated on data collected in the wild before its results can be accepted in practice. Secondly, the model in its current form does not take into account any type of lag effect from transferring the insects from one temperature to another. This type of effect could be significant in our data, as the environments of some insects were warmed or cooled by up to 15 °C instantaneously. While not optimal, these temperature transfers are necessary to capture observations at temperatures near the upper and lower limits of the developmental range without excessive mortality in the data. Future efforts could take these changes into account by modelling the effect of temperature changes on the insects. The current state of the model also does not include any correlation in development rates within individuals, among developmental stages. Also, since we have individual-level data, in theory we could follow individuals across their development stages and use that shared information to get better estimates within the interval censored windows. Another area of further study is the use of continuous measurements such as weight and length, instead of discrete observations of moulting, to track development. Using these continuous measurements would eliminate the difficulties that come with fitting a model to interval censored data, since measurements could be interpolated between unobserved states. Body size is also a better predictor of fitness traits such as fecundity (Honěk 1993). However, these measurements are much more time-consuming to collect. In terms of the selection of a TPC to fit to our data, the Schoolfield (1981) model provides an advantage over TPCs with sharp temperature thresholds, specifically for the Bayesian implementation of this model. With our implementation of individual variation, each individual would have exactly the same temperature thresholds for development, but we do not expect this to be the case. Using a development rate curve that gradually approaches zero provides more flexibility. Furthermore, when dealing with the gradient of a likelihood surface, blunt thresholds create pathologies and therefore are not a desirable characteristic of a TPC in our framework. An example of a TPC with upper and lower temperature thresholds is the curve presented in Regniere and St. Amant (2012).

4.2. POSTERIOR SAMPLES

While the typical MCMC diagnostics of $\hat{R}$, bulk and tail ESS, divergent transitions and posterior predictive checks did not indicate model misfit, lack of convergence or other difficulties with MCMC sampling, the SBC did reveal some issues in the model fitting. The non-uniformity in the SBC histograms indicates that the sampler does not properly recover the parameters for the scale of individual variation in the double interval-censored stages, nor the development rates at the transfer temperatures in the double interval-censored

stages. These patterns in the histograms suggest that both the interval censoring and the temperature transfers intrinsic to the data make it difficult to fit. The boxplots from the SBC of the σ_ϵ parameters show discrepancies in the location, scale and skew of the prior and posterior distributions. To use the model fit in practice, we suggest applying corrections to the posterior distribution according to the mismatch found in the SBC. For example, we can apply a Box–Cox transformation to the SBC-derived posterior distribution to get it as close to normal as we can, and log transform the log-normally distributed prior samples to make them normal as well. We can then use robust measures of location and scale (median and scaled IQR) to obtain a Gaussian distribution and apply the mean and standard deviation of the log prior to have the distributions match. Using the power parameter from the initial Box–Cox transformation, the location and scale parameters derived from the log prior and the SBC-derived posterior, one could then transform the data-derived posterior in a similar fashion. Performing this transformation improved the 90% coverage of the posteriors in the SBC for the σ_ϵ parameters, but the resulting credible intervals still undercover, and thus, these results should be interpreted with caution.

4.3. SIMULATIONS

We can obtain credible intervals for values of interest by sampling populations and individuals directly from the posterior distribution and using these values for simulations based on realistic temperature profiles (Fig. 5). Current simulations using weather data to predict spruce budworm population development rely on point estimates of model parameters, with the only stochastic component coming from individual-level variation around a deterministic development curve (Régnière et al. 2014; Régnière et al. 2012). This simulation method neglects uncertainty around parameter estimates and therefore underestimates the uncertainty in downstream predictions. In contrast, using Bayesian methods to sample both individual- and population-level variation from the estimated posterior distributions propagates the full range of model uncertainty through the simulations, so that they are reflected in the credible intervals for predicted values. This technique is equally applicable to scenarios informed by current conditions, based on historical weather data, or using climate projections to predict how insect populations will develop in the future.

[Received March 2022. Revised October 2023. Accepted October 2023.]

REFERENCES

Balakrishnan N, Leung M (1988) Order statistics from the type I generalized logistic distribution. Commun Stat Simul Comput 17(1):25–50

Bean J (1961) Predicting emergence of second-instar spruce budworm larvae from hibernation under field conditions in Minnesota. Ann Entomol Soc Am 54(2):175–177

Bewick S, Cantrell RS, Cosner C, Fagan WF (2016) How resource phenology affects consumer population dynamics. Am Nat 187(2):151–166

Blais JR (1983) Trends in the frequency, extent, and severity of spruce budworm outbreaks in eastern Canada. Can J For Res 13(4):539–547

Bolker BM (2008) Ecological models and data in R. Princeton University Press

Calle ML (2002) The Analysis of interval censoring and double censoring via Markov chain Monte Carlo methods. Technical report, Universitat de Vic—Central University of Catalonia

Cameron D, McDougall G, Bennett C (1968) Relation of spruce budworm development and balsam fir shoot growth to heat units. J Econ Entomol 61(3):857–858

Candau J-N, Dedes J, Lovelace A, MacQuarrie C, Perrault K, Roe A, Studens K, Wardlaw A (2019) Validation of a spruce budworm phenology model across environmental and genetic gradients: applications for budworm control and climate change predictions. Technical report, SERG International

Childress ES, Letcher BH (2017) Estimating thermal performance curves from repeated field observations. Ecology 98(5):1377–1387

Chuine I, Régnière J (2017) Process-based models of phenology for plants and animals. Annu Rev Ecol Evol Syst 48:159–182

Cook SR, Gelman A, Rubin DB (2006) Validation of software for bayesian models using posterior quantiles. J Comput Graph Stat 15(3):675–692

Corkrey R, Olley J, Ratkowsky D, McMeekin T, Ross T (2012) Universality of thermodynamic constants governing biological growth rates. PLoS ONE 7(2):e32003

Crimmins TM, Gerst KL, Huerta DG, Marsh RL, Posthumus EE, Rosemartin AH, Switzer J, Weltzin JF, Coop L, Dietschler N et al (2020) Short-term forecasts of insect phenology inform pest management. Ann Entomol Soc Am 113(2):139–148

Damos P, Savopoulou-Soultani M (2012) Temperature-driven models for insect development and vital thermal requirements. Psyche, 2012. Article ID 123405

Dennis B, Kemp W, Beckwith R (1986) Stochastic model of insect phenology: estimation and testing. Environ Entomol 15(3):540–546

Gajewski Z, Stevenson LA, Pike DA, Roznik EA, Alford RA, Johnson LR (2021) Predicting the growth of the amphibian chytrid fungus in varying temperature environments. Ecol Evol 11(24):17920–17931

Gelman A, Vehtari A, Simpson D, Margossian CC, Carpenter B, Yao Y, Kennedy L, Gabry J, Bürkner P-C, Modrák M (2020) Bayesian workflow. arXiv preprint arXiv:2011.01808

Hoffmann AA, Chown SL, Clusella-Trullas S (2013) Upper thermal limits in terrestrial ectotherms: how constrained are they? Funct Ecol 27(4):934–949

Honěk A (1993) Intraspecific variation in body size and fecundity in insects: a general relationship. Oikos 66:483–492

Hudes E, Shoemaker C (1988) Inferential method for modeling insect phenology and its application to the spruce budworm (Lepidoptera: Tortricidae). Environ Entomol 17(1):97–108

Ikemoto T (2005) Intrinsic optimum temperature for development of insects and mites. Environ Entomol 34(6):1377–1387

Johnson LR, Ben-Horin T, Lafferty KD, McNally A, Mordecai E, Paaijmans KP, Pawar S, Ryan SJ (2015) Understanding uncertainty in temperature effects on vector-borne disease: a Bayesian approach. Ecology 96(1):203–213

Kristensen K, Nielsen A, Berg CW, Skaug H, Bell B (2015) TMB: automatic differentiation and Laplace approximation. arXiv preprint arXiv:1509.00660

Lawrence RK, Mattson WJ, Haack RA (1997) White spruce and the spruce budworm: defining the phenological window of susceptibility. Can Entomol 129(2):291–318

LaZerte SE, Albers S (2018) weathercan: download and format weather data from environment and climate change Canada. J Open Source Software 3(22):571

Lysyk T (1989) Stochastic model of eastern spruce budworm (Lepidoptera: Tortricidae) phenology on white spruce and balsam fir. J Econ Entomol 82(4):1161–1168

McCarthy MA, Masters P (2005) Profiting from prior information in bayesian analyses of ecological data. J Appl Ecol 42:1012–1019

McManis AE, Powell JA, Bentz BJ (2018) Developmental parameters of a southern mountain pine beetle (Coleoptera: Curculionidae) population reveal potential source of latitudinal differences in generation time. Can Entomol 151:1

Monnahan C, Kristensen K (2018) No-U-Turn sampling for fast Bayesian inference in ADMB and TMB: introducing the adnuts and tmbstan R packages. PloS ONE 13(5):e0197954

Mordecai EA, Caldwell JM, Grossman MK, Lippi CA, Johnson LR, Neira M, Rohr JR, Ryan SJ, Savage A, Shocket MS et al (2019) Thermal biology of mosquito-borne disease. Ecol Lett 22(10):1690–1708

Neal RM (1993) Probabilistic inference using Markov chain Monte Carlo methods. University of Toronto Toronto, ON, Canada, Department of Computer Science

Perrault K, Wardlaw A, Candau J-N, Irwin C, Demidovich M, MacQuarrie C, Roe A (2021) From branch to bench: establishing wild spruce budworm populations into laboratory colonies for the exploration of local adaptation and plasticity. Can Entomol 153(3):374–390

Porter J, Parry M, Carter T (1991) The potential effects of climatic change on agricultural insect pests. Agric For Meteorol 57(1–3):221–240

Pruess KP (1983) Day-degree methods for pest management. Environ Entomol 12(3):613–619

Quinn BK (2017) A critical review of the use and performance of different function types for modeling temperature-dependent development of arthropod larvae. J Therm Biol 63:65–77

Rebaudo F, Rabhi V-B (2018) Modeling temperature-dependent development rate and phenology in insects: review of major developments, challenges, and future directions. Entomol Exp Appl 166(8):607–617

Régnière J (1984) A method of describing and using variability in development rates for the simulation of insect phenology. Can Entomol 116(10):1367–1376

Régnière J, St-Amant R, Duval P (2012) Predicting insect distributions under climate change from physiological responses: spruce budworm as an example. Biol Invasions 14(8):1571–1586

Régnière J (1987) Temperature-dependent development of eggs and larvae of Choristoneura fumiferana (Clem.) (Lepidoptera: Tortricidae) and simulation of its seasonal history. Can Entomol 119(7–8):717–728

Régnière J, Powell J, Bentz B, Nealis V (2012) Effects of temperature on development, survival and reproduction of insects: experimental design, data analysis and modeling. J Insect Physiol 58(5):634–647

Régnière J, Saint-Amant R, Béchard A, Moutaoufik A (2014) BioSIM 10: User's manual. Laurentian Forestry Centre

Schoolfield R, Sharpe P, Magnuson C (1981) Non-linear regression of biological temperature-dependent rate models based on absolute reaction-rate theory. J Theor Biol 88(4):719–731

Shocket MS, Verwillow AB, Numazu MG, Slamani H, Cohen JM, El Moustaid F, Rohr J, Johnson LR, Mordecai EA (2020) Transmission of West Nile and five other temperate mosquito-borne viruses peaks at temperatures between 23C and 26C. eLife 9:e58511

Stan Development Team (2020) Stan modeling language users guide and reference manual 2.25

Stedinger J, Shoemaker C, Tenga R (1985) A stochastic model of insect phenology for a population with spatially variable development rates. Biometrics 41(3):691–701

Sturtevant BR, Achtemeier GL, Charney JJ, Anderson DP, Cooke BJ, Townsend PA (2013) Long-distance dispersal of spruce budworm (*Choristoneura fumiferana* Clemens) in Minnesota (USA) and Ontario (Canada) via the atmospheric pathway. Agric For Meteorol 168:186–200

Talts S, Betancourt M, Simpson D, Vehtari A, Gelman A (2018) Validating Bayesian inference algorithms with simulation-based calibration. arXiv preprint arXiv:1804.06788

Uvarov B (1931) Insects and climate. Trans R Entomol Soc London 79:1

Vehtari A, Gelman A, Simpson D, Carpenter B, Bürkner P-C (2021) Rank-normalization, folding, and localization: an improved $\widehat{R}$ for assessing convergence of MCMC. Bayesian Anal 16(2):667–718

Villena OC, Ryan SJ, Murdock CC, Johnson LR (2022) Temperature impacts the environmental suitability for malaria transmission by Anopheles gambiae and Anopheles stephensi. Ecology 103(8):e3685

von Schmalensee L, Hulda Gunnarsdóttir K, Näslund J, Gotthard K, Lehmann P (2021) Thermal performance under constant temperatures can accurately predict insect development times across naturally variable microclimates. Ecol Lett 24(8):1633–1645

Wardlaw A, Perrault K, Roe A, Dedes J, Irwin C, MacQuarrie C, Candau J-N (2022) Methods for estimating and modelling spruce budworm development rates at constant temperatures. Can Entomol 154:e9