Supporting Text: Power Spectra Reveal the Influence of Stochasticity on Nonlinear Population Dynamics

Daniel C. Reuman, Robert A. Desharnais, Robert F. Costantino, Omar S. Ahmad, Joel E. Cohen

Contents

1: Methods

2: Supporting Results

1: Methods

1.1: Models and parameter values

Fig. 5: Models of *T. castaneum* population dynamics. The deterministic Larvae-Pupae-Adult model (LPA) is presented in the upper left corner, box [1]. The stochastic demographic Larvae-Pupae-Adult model (SD-LPA) is displayed in the upper right corner, box [2]. The lattice stochastic demographic Larvae-Pupae-Adult model (LSD-LPA) is shown in the lower right corner, box [3]. We consider the LSD-LPA model to be assembled starting with the LPA model, adding stochasticity first, then constraining to integer populations, rather than *vice versa*. $[x] = \max(x,0)$. The trivariate random variables $E_t = (E_{1t}, E_{2t}, E_{3t})$ were standard normal with covariance matrix Σ, and were independent for different values of t. *See separate file for the figure.*

Table 1: The parameters used for the models of Fig. 5. These were obtained using time domain fitting methods (Dennis *et al.* 2001, Cushing *et al.* 2003). *See separate file for the table.*

1.2: Power spectrum fit tests

The "spectrum distance fit test" quantifies the frequency-domain fit between a model with specified parameters and an experimental time series, $\{x_0, \ldots, x_{T-1}\}$. Steps for computing the p-value result of the test for a population model and data with life stages L, P and A are:

- Use the model with initial conditions x_0 to produce N simulated time series of length T.
- Compute the log spectral estimate $\log(f_X^{(T)}(\lambda))$ of each life stage of each simulated time series ($3*N$ log spectra) using a "consistent spectral estimate" of Brillinger (Brillinger 2001, p. 146). The variable λ denotes normalized frequency (n.f.).
- For each $\lambda = 2s/T$ (for s an integer from 0 to $\lfloor T/2 \rfloor$), and for each life-stage, compute the mean of the N values of $\log(f_X^{(T)}(\lambda))$ that were computed for that life stage. Call the resulting three spectra (one for each life-stage) the "mean model log spectra," $\bar{u}_L^{(T)}(\lambda)$, $\bar{u}_P^{(T)}(\lambda)$ and $\bar{u}_A^{(T)}(\lambda)$.
- For each log spectrum of a model-generated time series (the three life-stage components of which we denote $u_L^{(T)}(\lambda)$, $u_P^{(T)}(\lambda)$ and $u_A^{(T)}(\lambda)$) compute the three distances:

  Eqn. 1
  $$d_w(u_w,\bar{u}_w) = \sqrt{\sum_{s=0}^{\lfloor T/2 \rfloor} (u_w^{(T)}(2s/T) - \bar{u}_w^{(T)}(2s/T))^2}$$

  for $w = L, P, A$. Let
  $$d(u,\bar{u}) = \sqrt{\sum_{w=L,P,A} d_w(u_w,\bar{u}_w)^2} \ .$$

  This $d(u, \bar{u})$ gives N points from a distribution of d.
- Denote log-spectra of the life-stage components of the experimental data time series by $e_L^{(T)}(\lambda)$, $e_P^{(T)}(\lambda)$, $e_A^{(T)}(\lambda)$, and compute $d(e, \bar{u})$.

- Generate an approximate *p*-value from the *d* distribution for this experimental replicate by computing the percentage of the $d(u, \bar{u})$ that are greater than $d(e, \bar{u})$.

The distance in Eqn. 1 is the square root of the sum of the squared differences between the mean model spectrum values $\bar{u}_w(2s/T)$ and the individual spectrum values $u_w(2s/T)$. This quantity is small when the spectra u_w and $\bar{u}_w$ are close for all values of s simultaneously.

Steps for computing the *p*-value result of the "spectrum shape fit test" for a population model and data with life stages L, P and A are:

- Use the model with initial conditions x_0 to produce N simulated time series of length T.
- Compute the log spectral estimate $\log(f_X^{(T)}(\lambda))$ of each life stage of each simulated time series ($3*N$ log spectra) using a "consistent spectral estimate" of Brillinger (Brillinger 2001, p. 146). The variable λ denotes normalized frequency (n.f.).
- For each $\lambda = 2s/T$ (for s an integer from 0 to $\lfloor T/2 \rfloor$), and for each life-stage, compute the mean of the N values of $\log(f_X^{(T)}(\lambda))$ that were computed for that life stage. Call the resulting three spectra (one for each life-stage) the "mean model log spectra," $\bar{u}_L^{(T)}(\lambda)$, $\bar{u}_P^{(T)}(\lambda)$ and $\bar{u}_A^{(T)}(\lambda)$.
- For each log spectrum of a model-generated time series (the three life-stage components of which we denote $u_L^{(T)}(\lambda)$, $u_P^{(T)}(\lambda)$ and $u_A^{(T)}(\lambda)$) compute the three correlation coefficients $c(u_w^{(T)}(\lambda), \bar{u}_w^{(T)}(\lambda))$ for $w = L, P, A$, and compute their sum, $c(u, \bar{u})$. This gives N points from a distribution of c values.
- Denote log-spectra of the life-stage components of the experimental data time series by $e_L^{(T)}(\lambda)$, $e_P^{(T)}(\lambda)$, $e_A^{(T)}(\lambda)$, and compute $c(e, \bar{u})$.
- Generate an approximate *p*-value from the c distribution for this experimental replicate by computing the percentage of the $c(u, \bar{u})$ that are less than $c(e, \bar{u})$.

These tests give approximate *p*-values because they are based on stochastic simulation: executing the same steps twice gives slightly different results. For this study, every *p*-value was computed 100 times, and means and standard deviations of the resulting distribution of approximate *p*-values were computed, as well as minima, 25th, 50th, and 75th percentiles, and maxima. The variance of the *p*-value distributions obtained from the tests can be decreased by increasing N. Typical values of N used in this study were 5000 and 10,000.

A model for which the 25th percentile of a *p*-value distribution from one of the two tests above was bigger than 1% was considered to be well-fitting with respect to that test. Models that fit for both tests using this criterion produced time series with power spectra that looked visually very similar to those of the data time series (e.g., Fig. 2 A, B, C, E but not D, F).

1.3: Non-equivalency of fit in the time and frequency domains

A model that had good one-time-step forecasting ability did not always have to fit well in the frequency domain according to the spectrum distance fit test, nor conversely. A stochastic version of the Ricker model,

$$p_{n+1} = 17 * p_n \exp(-\frac{1}{300} p_n + 0.03 * X_n),$$

was used to generate three data time series d_1, d_2, d_3 of length 128 starting from initial condition $p_0 = 14$. The X_n were independent standard normal. One-step-forecasting maximum-likelihood techniques similar to those used in Cushing *et al.* (2003) were used to fit the model

$$p_{n+1} = 17 * p_n \exp(-\frac{c}{300} p_n + 0.13 * X_n)$$

with d_1, d_2, d_3 (the coefficient 0.13 was intentionally different from the coefficient 0.03 in the above model); the best-fitting model had $c \approx 1$. The spectrum distance fit test found best-fitting $c \approx 1.24$. In the frequency domain, the model with $c = 1$ had $p = 0.02\%$ (spectrum distance fit test); the model with $c = 1.24$ had $p = 7.76\%$ (spectrum distance fit test). P-values of three replicates were combined into a single p-value by the Fisher method. The time-domain likelihood function evaluated at $c = 1$ was about 1.7×10^{181}; evaluated at $c = 1.24$ it was less than 10^{-10}. Thus the time-domain fitting indicates that $c = 1.24$ is much less likely than $c = 1$, while the spectrum distance fit test indicates the reverse.

A second-order-Markov relative of the stochastic Ricker model was used to generate three data time series e_1, e_2, e_3 of length 128 starting from initial conditions $p_0 = 14$, $p_1 = 227.1485$:

$$p_{n+1} = (17 * p_n + 0.1 * p_{n-1}) \exp(-\frac{1}{300} p_n + 0.03 * X_n).$$

The dependence of p_{n+1} on p_{n-1} can be interpreted as some eggs laid during time step n-1 not hatching until time step n+1, although most hatch during time step n. The parameters 17 and 0.1 in this model correspond to the biological hypothesis that 17/(17+0.1) = 99.42% of eggs hatch after one time step, while the rest hatch after two time steps. One-step-forecasting maximum-likelihood techniques similar to those used in Cushing *et al.* (2003) were used to fit the model

$$p_{n+1} = b * p_n \exp(-\frac{1}{300} p_n + \sigma * X_n)$$

with e_1, e_2, e_3. This model corresponds to the reasonable but inaccurate biological hypothesis that all eggs hatch after one time step. The best-fitting model had $b \approx 17.62$ and $\sigma \approx 0.0688$. The spectrum distance fit test found best-fitting $b \approx 16.60$ and $\sigma \approx 0.0113$; these parameters gave $p = 15.24\%$ in the frequency domain (spectrum distance fit test). The parameters $b \approx 17.62$ and $\sigma \approx 0.0688$ gave $p = 0.56\%$ in the frequency domain (spectrum distance fit test). P-values of three replicates were combined into a single p-value by the Fisher method. The time-domain likelihood function evaluated at $(b, \sigma) = (17.62, 0.0688)$ was about 2.3×10^{208}; evaluated at $(b, \sigma) = (16.60, 0.0113)$ it was less than 10^{-10}.

To show that the methods of this study are capable of rejecting reasonable model functional forms based on data of 41-time-step length, we used the model

$$p_{n+1} = (17 * p_n + 0.25 * p_{n-1}) \exp(-\frac{1}{300} p_n + 0.03 * X_n)$$

to generate time series f_1, f_2, f_3 of length 41 starting from the initial conditions $p_0 = 14$, $p_1 = 227.1485$. The parameters 17 and 0.25 in this model correspond to the biological hypothesis that 17/(17+0.25) = 98.55% of eggs hatch after one time step, while the rest hatch after two time steps. One-step-forecasting maximum-likelihood techniques were used to fit the model

$$p_{n+1} = b * p_n \exp(-\frac{1}{300} p_n + \sigma * X_n)$$

with f_1, f_2, f_3. The best-fitting model had $b \approx 18.84$ and $\sigma \approx 0.1630$; the frequency domain fit of the model with these parameters gave $p < 0.0001$ (spectrum distance fit test; the 3 p-values were combined via the Fisher method). The spectrum distance fit test found best-fitting $b \approx 16.00$, $\sigma \approx 0.0767$, with p = 2.44%. The time-domain likelihood function evaluated at $(b, \sigma) = (18.84, 0.1630)$ was about 5.0×10^{20}; evaluated at $(b, \sigma) = (16.00, 0.0767)$ it was less than 10^{-10}. No single set of parameters (b, σ) produced a fit according to both time and frequency domain measures of fit, so we reject the model.

We used the model

$$p_{n+1} = (17 * p_n + 0.5 * p_{n-1}) \exp(-\frac{1}{300} p_n + 0.03 * X_n)$$

to generate time series f_1, f_2, f_3 of length 41 starting from the initial conditions $p_0 = 14$, $p_1 = 227.1485$. The parameters 17 and 0.5 in this model correspond to the biological hypothesis that 17/(17+0.5) = 97.14% of eggs hatch after one time step, while the rest hatch after two time steps. One-step-forecasting maximum-likelihood techniques were used to fit the model

$$p_{n+1} = b * p_n \exp(-\frac{1}{300} p_n + \sigma * X_n)$$

with f_1, f_2, f_3. The best-fitting model had $b \approx 22.40$ and $\sigma \approx 0.3850$; the frequency domain fit of the model with these parameters gave $p < 0.0001$ (spectrum distance fit test; 3 p-values combined by the Fisher method). The spectrum distance fit test found best-fitting $b \approx 16.33$, $\sigma \approx 0.08$, with $p = 0.0009$. Frequency domain fit was not possible with this model functional form for any parameters in the range searched.

1.4: Lag metric and relative lag metric

Two “lag metrics” were used: one developed previously (King *et al.* 2004), and an adaptation of that method. The adapted method, called the “relative lag metric,” detected the presence of length-n approximately-periodic patterns in a time series. For the time series $x_0, \ldots, x_{T-1}$, the “relative lag metric” method computed

$$D_t = (\sum_{i=t-n+1}^{t} \delta_1(x_i, x_{i-n}))/n$$

for $t = 2n-1$ to $T-1$. The notation δ_1 represented distance under the 1-norm:

$$\delta_1((L_i, P_i, A_i), (L_j, P_j, A_j)) = | L_i - L_j | + | P_i - P_j | + | A_i - A_j |.$$

The quantity D_t was the average distance between x_i and x_{i-n} for i ranging over the past n time steps; low values of D_t indicated that the data had been passing through an approximately repeating length-n periodic pattern for at least the last $2n$ time steps (one sequence of n points followed by one approximate repetition of that sequence). The distance D_t was plotted versus t, keeping track of the times for which it was less than 55 (the value 55 was taken from a previous application of the lag metric to the *T. castaneum* system, King *et al.* 2004). The number of total t values for which D_t was less than 55 and the maximum number of contiguous D_t values below 55 provided an indication of the degree to which approximately-repeating length-n patterns occurred.

The original "lag metric" method (King *et al.* 2004) differed from the relative lag metric because it took a fixed length-n pattern $c^{(n)}$ as additional input. It measured the average distance between each n-long sub-time series $x_{t-n+1}, \ldots, x_t$ and $c^{(n)}$, and plotted this distance against t. The pattern $c^{(n)}$ had n possible rotations; the lag metric plotted the average distance to each possible rotation, giving n different lines on a single plot. When any one of these lines was below 55 the data were considered to be approximately repeating the pattern $c^{(n)}$.

1.5: Extraction and significance of length-7 pattern

This section gives the details of how an approximately-repeating pattern of length 7 was extracted from LSD-LPA model output and from experimental data with $c_{pa} = 0$, and how the statistical significance of the presence of that pattern was determined. An approximately repeating pattern of length 6 was extracted from model output and experimental data with $c_{pa} = 0.5$ using the same methods.

Extraction:

Starting with length-1024 output of the LSD-LPA model with $c_{pa} = 0$ and $\mu_a = 0.96$, and other parameters equal to those of Table 1, all time steps for which the relative lag metric, with lag 7 (Section 1.4), was less than 75 were found. Of these, the time steps were kept that were part of a run of length at least 8 time steps with relative lag metric below 75. Such a run indicated that the time series was executing an approximately-periodic pattern of length 7, for at least 21 time steps (3 repetitions of the pattern). For the model output used, there were no runs of length greater than 10 time steps for which the relative lag metric was below 75.

If t was the first time step of each run of length 8 or greater, the population triples from t-13 to t+7 were extracted, giving 3 potential repetitions of a length-7 pattern. In this way, 27 potential patterns of length 7 were extracted from the length-1024 series. The first of these (an arbitrary choice) was labeled x. For y equal to each of the 26 other patterns, the average 1-norm distance between y and x was measured for each of the 7 possible rotations of y. The patterns y for which no rotation had distance to x less than 90 were discarded. No pattern y had more than one rotation with distance to x less than 90. This one rotation "lined up" with x. After discarding the patterns for which no rotation lined up, the adult life stage of 10 of the 19 remaining patterns was plotted, randomly chosen (Fig. 3D), and the median of the 19 population values for each time step and for each life stage was computed to get one median length-7 pattern.

The (original) lag metric (Section 1.4) was used to compare the experimental time series from the treatment with $c_{pa} = 0$ (replicates 5, 12 and 15) to the median length-7 pattern extracted from model output. For replicates 5 and 15, the metric never went below the threshold of 55, but for replicate 12 it did so for 4 consecutive time steps, indicating that this length-7 pattern occurred in experimental data. The minimum lag metric value for replicate 12 was 40.14 individuals, and for replicate 5 it was 62.57 individuals. The 7 time steps ending at the time step with minimal lag metric were extracted from replicates 5 and 12, and rotated to line up with the length-7 patterns extracted from typical model output. The adult life stages of these extracted patterns were plotted (Fig. 3D).

Statistical Significance:

To test whether the median pattern extracted and displayed in Fig. 3D (open circles) occurred only by chance, we randomly permuted, in 1000 different ways, the same LSD-LPA model output from which the pattern was extracted. For each permuted time series, the same steps were undertaken as were used to extract the pattern from the un-permuted model output: the relative lag metric (Section 1.4) was applied to each permuted time series, and the length of the longest run of time steps for which the relative lag metric was below the threshold 75 was recorded. None of the 1000 permutations had a run longer than 6. Nine runs of length 8 or more occurred in the un-permuted model output.

To test whether the median pattern displayed in Fig. 3D (open circles) was detected in experimental replicates 5 and 12 by chance alone, we randomly permuted, in 1000 different ways, the experimental time series of replicates 5 and 12, and attempted to find the same median pattern in each permuted time series using the same methods with which it was detected in un-permuted experimental data: the lag metric was used to test for the presence of the median pattern in each permuted time series. The lag metric never went below the threshold of 55 for any of the 1000 permutations for either replicate. The same test was run on data from experimental replicate 15 (same treatment group), with the same results.

1.6: Plotting spectral peaks versus Σ factor (Fig. 4)

For each of 101 log Σ factors evenly spaced from $\log(2 \times 10^{-5})$ to 1, 1000 time series of length $T = 1024$ were generated using the SD-LPA model. The log periodogram spectral estimate (Brillinger 2001) of each life stage of each time series was computed. For each n.f. value of the form $2s/T$ for s an integer from 1 to $\lfloor T/2 \rfloor$, the median of the log periodogram values at that n.f. was computed for each life stage. The periodogram value at n.f. = 0 was removed before computing medians because it did not contain spectral information.

The resulting median log periodograms (with the n.f. = 0 value omitted) were smoothed using local averaging to eliminate small spurious peaks and valleys. The smoothing was simple averaging of the median log-periodogram value at each frequency with the values at the 10 frequencies of the form $2s/T$ (for s an integer from 1 to $\lfloor T/2 \rfloor$) on each side of it. For most n.f. values, 21 quantities were averaged. The average of all available frequencies was used for normalized frequencies close to 0 or 1 not having ten frequencies on one side.

Locations of local maxima and minima were computed using the smoothed median log periodogram for each log Σ factor. A frequency was considered to be the location of a local maximum if its smoothed median log-periodogram value was at least as great as the values of the 10 frequencies on each side of it (or fewer for frequencies close to 0 and 1). Heights of the local maxima and minima were computed using the smoothed median log periodogram. Locations and heights of spectral peaks and valleys were plotted against log Σ factor (Fig. 4).

Although not all trends of peaks appearing on plots produced in this way will be biologically meaningful (some are artifacts of the spectrum method), they will all be highly statistically significant for two reasons. First, computing the median of 1000 logperiodigrams and subsequently smoothing the median eliminates peaks that are due to dramatic, chance departures of dynamics from central patterns. Second, the random components affecting simulations with adjacent log Σ factors are independent. Although a statistically rare and not dynamically meaningful peak may occasionally occur for a single value of the log Σ factor, the probability that an entire trend of peaks will materialize by chance is vanishingly small.

1.7: Linearization theory

Theory based on linearization uses the power spectrum to explain interactions between stochasticity and nonlinear dynamics in physical systems (Wiesenfeld 1985) and population models (Greenman & Benton 2005). We briefly review the main qualitative conclusions of the theory. Given a deterministic model $\mathbf{x}_{t+1} = f(\mathbf{x}_t)$ with a finite stable attractor $A = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n)$ consisting of n points periodically repeated, the spectral effects of adding normally-distributed stochasticity $\boldsymbol{\varepsilon}$ to the model, $\mathbf{x}_{t+1} = f(\mathbf{x}_t) + \boldsymbol{\varepsilon}$, can be predicted from the eigenvalues and eigenvectors of the "susceptibility matrix" $S = J_{\mathbf{x}_n} J_{\mathbf{x}_{n-1}} \ldots J_{\mathbf{x}_1}$, where $J_{\mathbf{x}}$ is the Jacobian of f at $\mathbf{x}$. A negative real eigenvalue d_1 with corresponding eigenvector $\mathbf{v}$ indicates that small perturbations $\mathbf{h}$ from $\mathbf{x}_1$ in the direction of $\mathbf{v}$ will become, in n time steps, approximately $S\mathbf{h} = d_1\mathbf{h}$, which has direction opposite that of $\mathbf{h}$. This direction switching of perturbations has period $2n$, twice that of the deterministic system. It is over-compensatory decay of the perturbation $\mathbf{h}$. If $\boldsymbol{\varepsilon}$ has a component in the direction of $\mathbf{v}$, and $\mathbf{v}$ has non-zero i^{th} component, then direction switching of perturbations (over-compensatory decay) causes a stochastically-induced peak in the spectrum of the i^{th}-component time series at n.f. $1/n$ (half the frequency of the deterministic system if $n > 1$, because the deterministic system has n.f. $2/n$).

If S has a complex-conjugate pair of eigenvalues d_1 and d_2 with phase angles α and $-\alpha$, then there is a 2-dimensional space $\mathbf{V}$ such that S approximately rotates vectors in $\mathbf{V}$ by average angle α and then re-scales them by $|d_1|$. Perturbations $\mathbf{h}$ from $\mathbf{x}_1$ that lie in $\mathbf{V}$ will be approximately $S\mathbf{h}$ in n time steps: they will be approximately rotated by α and multiplied by $|d_1|$. This amounts to oscillatory decay of the perturbation $\mathbf{h}$, with oscillations of n.f. $\alpha/\pi n$. If stochasticity produces perturbations with components in $\mathbf{V}$, and if $\mathbf{V}$ contains vectors with non-zero i^{th} component, then rotation causes a stochastically-induced peak at n.f. $\alpha/\pi n$ in the spectrum of the i^{th}-component time series. Because rotation by α is the same as rotation by $2\pi - \alpha$ in the other direction, a second peak at $(2\pi-\alpha)/\pi n$ also occurs if $n > 1$.

If S has a positive real eigenvalue d_1 with corresponding eigenvector $\mathbf{v}$, then small perturbations $\mathbf{h}$ from $\mathbf{x}_1$ in the direction of $\mathbf{v}$ will become approximately $d_1\mathbf{h}$ in n time steps. By the assumed stability, $d_1 < 1$, so perturbations $\mathbf{h}$ decay exponentially over time, with larger d_1 causing slower decay. The decay is under-compensatory. If $\boldsymbol{\varepsilon}$ has a component in the direction of $\mathbf{v}$ and $\mathbf{v}$ has non-zero i^{th} component, then under-compensatory decay of perturbations produces reddening in the i^{th}-component time series: spectral power increases at low frequencies.

Because A is stable, all eigenvalues have magnitude less than 1. Larger eigenvalue magnitudes cause larger stochastically-induced peaks for any combination of eigenvalues. Larger stochasticity also causes larger peaks, up to the limit where linearization theory begins breaking down. Linearization theory also makes quantitative spectral predictions (Wiesenfeld 1985; Greenman & Benton 2005); this study makes use of qualitative predictions only.

1.8: Computations

All logarithms were base 10. The Matlab function ‘fft’ was used for discrete Fourier transforms. All computations were done on a 3.2 GHz desktop PC using Matlab versions 6.5.0 and 6.5.1, with one exception. An adaptation of the Nelder-Mead simplex algorithm was used to optimize the LSD-LPA model fit starting from parameters of Table 1. This algorithm was run on a cluster of 16 PC machines of 1.5 GHz each. The Matlab C compiler was used to translate code to the C programming language for this purpose. The algorithm was stochastic, so it was repeated many times. The computational power of the cluster of 16 machines was needed for this repetition.

2: Supporting Results

2.1: Frequency-domain fit and predictive ability of the LSD-LPA model

Using model parameters obtained by maximum-likelihood time-domain fitting (Section 1.1), the frequency-domain fit of the LSD-LPA model with all 24 data replicates was tested (Fig. 1, box 2). The six replicates of length 213 were truncated to length 41 for this purpose. The fit-testing algorithms were based on stochastic simulations, so multiple tests of fit with the same data and parameters produced slightly different *p*-values; but the standard deviation of *p*-values for 100 tests was never more than 1.25% for any replicate, for both the spectrum distance fit test and the spectrum shape fit test. All replicates but one for the spectrum distance fit test and two for the spectrum shape fit test had at least 75 of 100 approximate *p*-values above 1%. For all replicates and life stages, visual confirmation of fit was obtained (Fig. 2A). An adaptation of the Nelder-Mead simplex algorithm was used to make small modifications of parameters to decrease the L_2 distances between data log spectra and the mean of log spectra of model-generated time series. No significant improvement in frequency-domain fit occurred (according to the spectrum distance fit test). The model fitted better with the given parameters than with any nearby parameters; it fitted data well in both the time and frequency domains with the same parameters.

Spectral estimates from length-41 time series lack resolution. To make spectral predictions of higher resolution, we generated many time series of length 213 using the LSD-LPA model (parameter values were the same as those of Section 1.1); spectral estimates based on these time series were model-based hypotheses of population spectra (Fig. 1, box 3). We tested the predictive ability of the model and the strategy of this study (Fig. 1, fourth box) by comparing these hypotheses with spectral estimates from the 6 experimental replicates of length 213 (control and c_{pa}=0.35 replicates). Hypotheses and data-based estimates agreed qualitatively and quantitatively: a second spectral peak at normalized frequency (n.f.) 0.33 was predicted by the model and confirmed by the length-213 experimental time series for c_{pa}=0.35 (Fig. 2B); every replicate had at least 75 of 100 approximate-*p*-values above 1% for both spectrum fit tests. Standard deviations of *p*-value distributions were always less than 0.7%.

2.2: Fits of constrained models

We examined the fit in the time and frequency domains between several variants of the LSD-LPA model and the three experimental replicates with imposed $c_{pa} = 0.35$. Models used were: the constrained LSD-LPA model; and the constrained LSD-LPA model with additional constraints $c_{el} = 0$ and $c_{ea} = 0$, imposed separately. Parameters were computed that produced the best time-domain fit between each model and the three experimental replicates truncated to length 41 (Table 2). For all models, the time-domain-optimal parameters also gave a good fit with short time series in the frequency domain, according to both the spectrum distance fit test and the spectrum shape fit test (Table 3).

An adapted Nelder-Mead simplex algorithm was used to adjust model parameters to decrease the L_2 distance between log spectra of short experimental time series (length 41) and the mean of log spectra of many model-generated time series of the same length; the goal was to optimize frequency-domain fit. Because the algorithm optimized a stochastic objective function, it was run 5 times for each model, starting from initial parameters equal to the time-domain optimum. Results represent local minima of the mean of the stochastic objective function. Frequency-domain-optimized parameters often produced a slightly but not substantially better frequency-domain fit than did the time-domain-optimal parameters (Table 3); they sometimes fell outside the 99% confidence intervals of the time-domain-optimal parameters (Table 2).

Models with optimized parameters were used to generate many time series of length 213; the log spectra of these long time series were compared with log spectra of the three length-213 experimental time series with imposed $c_{pa} = 0.35$. P-values were produced describing the degree of similarity between model and experimental spectra (Table 3); these p-values reflect the ability of each model to make accurate predictions of spectra of long time series when optimized using short time series. Both time-domain-optimized and frequency-domain-optimized parameters were used, separately. The constrained LSD-LPA model was the only model capable of making accurate predictions, although all models fitted adequately with short time series in the time and frequency domains. This difference in predictive ability of the models arises because the constrained LSD-LPA model is the only model, among those considered here, that did not incorrectly omit a known biological mechanism.

The analytic strategy of this study (Fig. 1) combines the qualitative, biological information contained in a model functional form with quantitative information contained in short time series to produce the best available estimate of population spectra. A mechanistic model with the wrong mechanism will produce incorrect spectral predictions for long time series even though it may fit well with short time series.

Table 2: Optimized parameter values in the time domain (column 2), and in the frequency domain according to the spectrum distance fit test (columns 4-8), with respect to the length-41 experimental replicates with $c_{pa} = 0.35$ only. Values in brackets do not fall within the confidence intervals in column 3. Time-domain fitting used the maximum likelihood one-step forecasting methods of Dennis *et al.* (2001). Frequency-domain optimization used an adaptation of the Nelder-Mead simplex algorithm to minimize the distance between the mean log spectrum of model-generated time series and data log spectra. The time-domain-optimal parameters (column 2) were used as initial conditions. Confidence intervals for the time-domain-optimal parameters other than Σ were computed as in Dennis *et al.* (2001). Confidence intervals for entries of Σ were computed by re-sampling (with replacement) from the set of residuals of data from one-step forecasts of the model with time-domain-optimal parameters. *See separate file for the table.*

Table 3: Frequency domain fit of models with $c_{pa} = 0.35$ data, parameters as specified in Table 2. Table entries are p-values as given by the spectrum distance fit test (columns 2-3) and the spectrum shape fit test (columns 4-5). They describe the fit between models and length-41 (columns 2, 4) or length-213 (columns 3, 5) versions of the three $c_{pa} = 0.35$ experimental time series. *See separate file for the table.*

Figure 6: Frequency domain fit between the constrained LSD-LPA model (with no further constraint) and observed adult population time series of (A) length 41 and (B) length 213 from the three experimental replicates with $c_{pa} = 0.35$. The heavy dashed lines

are data log spectra. Light solid lines give the minimum, the 5th, 25th, 50th, 75th, and 95th percentiles, and the maximum values at each frequency value of 1000 log spectra of model-generated time series of (A) length 41 and (B) length 213. Triangles highlight the 5th and 95th percentiles. Parameters were optimized for length-41 data in the time domain (Table 2). Fit was similar when frequency-domain-optimized parameters were used (Table 3). Fit with other life stages was similar. Aliasing of fundamental (i.e., non-harmonic) frequencies of population fluctuation is unlikely to have occurred because the biology of *Tribolium* suggests that little fluctuation occurs for n.f. greater than 1. *See separate file for the figure.*

2.3: Analogues of Fig. 4 for other life stages and c_{pa} values

Figure 7: (A) Fig. 4A, reproduced here for easy comparison with the other panels. The analog of Fig. 4A for the pupal (B) and larval (C) life stages. (D) Fig. 4B, reproduced for comparison. The analogue of Fig. 4B for the pupal (E) and larval (F) life stages. *See separate file for the figure.*

Figure 8: Log power versus n.f. and log Σ factor, $c_{pa} = 0$, $\mu_a = 0.96$, for the adult (A), pupal (B), and larval (C) life stages. The smoothed median log periodogram estimate was used. As in Fig. 4B and Fig. 7, n.f. = 0 values have not been plotted. This is the 3-dimensional version of Fig. 4B. There appear to be differences among Figs. 7D-F. This figure shows that the differences are minor; they may not have biological significance. *See separate file for the figure.*

Figure 9: Analogues of Fig. 4 for other c_{pa} values. Adult life stage only. *See separate file for the figure.*

2.4: Limiting behavior of deterministic models for $c_{pa} = 0$

The limiting behavior of the LPA model is described by its attractor (Fig. 10). This attractor, an invariant loop, consists of two small circles: c_1 and c_2. All points on one of the circles have few larvae (33 to 39), many pupae (92 to 99), and few adults (30 to 35): a pupae-dominated age distribution. All points on the other circle have many larvae (115 to 124), few pupae (26 to 32) and many adults (93 to 100): a bimodal age distribution. Oscillations between heavily-pupae-dominated and heavily-bimodal distributions correspond to switching between opposite sides of c_1 and c_2; oscillations between moderately-pupae-dominated and moderately-bimodal distributions correspond to switching between the sides of c_1 and c_2 that are closest to each other.

Figure 10: The attractor (circles) and unstable equilibrium (x) of the LPA model for $c_{pa} = 0$, $\mu_a = 0.96$. *See separate file for the figure.*

2.5: Explanation of spectral changes with stochasticity

This section gives the details of: 1) the approximation made to apply linearization theory to the treatment with $c_{pa} = 0$; and 2) the experimental support for the proposed mechanisms behind the spectral changes depicted in Fig. 4.

Application of linearization theory to the $c_{pa}=0$ treatment:

The main text notes that linearization theory cannot immediately explain peak-change patterns for $c_{pa}=0$ because the square-root-scale LPA model with $c_{pa}=0$ has a non-finite attractor. However, the components of the attractor, $\sqrt{c_1}$ and $\sqrt{c_2}$, can substitute for a stable 2-point attractor. If $\mathbf{x}_t$ and $\mathbf{x}_{t+1}=f(\mathbf{x}_t)$ are square-root-scale population vectors in $\sqrt{c_1}$ and $\sqrt{c_2}$, respectively, then the product of Jacobians $S_{\mathbf{x}_t} = J_{\mathbf{x}_{t+1}} J_{\mathbf{x}_t}$ approximates a susceptibility matrix because $\mathbf{x}_{t+2}$ is close to $\mathbf{x}_t$. Using 5000 points $\mathbf{x}_t$ covering $\sqrt{c_1}$, the coefficient of variation of each matrix entry of $S_{\mathbf{x}_t}$ was never more than 7.9%. The matrices $S_{\mathbf{x}_t}$ all had similar eigenvalue and eigenvector structure. A negative eigenvalue $-0.12 \leq d_1 \leq -0.15$ always occurred (mean value -0.13); a complex-conjugate pair of eigenvalues $r\exp(\pm i\theta)$ with $0.97 \leq r \leq 1.03$ and $40.8° \leq \theta \leq 47.7°$ always occurred (mean θ was 44.4°, mean r was 0.999). The dominant eigenvalue was always complex, with mean phase angle 44.4°, so theory predicts a component of oscillatory decay of perturbations from the stable attractor of the square-root-scale model, and corresponding spectral peaks in the adult life stage at n.f. 44.4°/(180°*2)=0.123 and (360°-44.4°)/(180°*2)=0.877. The magnitude of the negative eigenvalue is small, and its corresponding eigenvector is almost perpendicular to the largest component of stochastic perturbations that affected the *Tribolium* system, so theory predicts that the corresponding stochastically-induced spectral peak will not be noticeable. Predicted stochastically-induced peaks are in the same place as deterministic model peaks. Linear theory predicts that locations of spectral peaks for $c_{pa}=0$ will not change with the addition of stochasticity.

For $c_{pa}=0$, linearization theory predicted changes in spectral peaks correctly for weak noise but incorrectly for stronger noise: the main (deterministic) LPA model peaks at 0.123, 0.887 and 1 were not greatly affected by weak demographic stochasticity (log Σ factors less than about -2.5), but were shifted unexpectedly for larger Σ factors (Fig. 4B). How can peak motion for higher Σ factors be explained? We answer this question in Fig. 11 below, and in the main text.

Experimental support for spectral changes:

We supported the theoretically-proposed mechanism producing the n.f. 0.33 spectral peak for $c_{pa} = 0.5$ with experimental data by examining perturbations of real population vectors from points in the LPA model attractor (Fig. 11A).

The LPA model with $c_{pa}=0$ had an unstable equilibrium, *e*, about midway between c_1 and c_2. The distance from any population vector to the attractor comprised of c_1 and c_2 was defined to be the minimum distance to any point in the attractor. For a range of Σ factors, means and 2.5th and 97.5th percentiles of this minimum distance were computed across all population vectors in a model-generated time series (Fig. 11B). Spectral peaks began to move for about the same Σ factors that 97.5th percentiles of distance distributions exceeded the distance between *e* and the attractor (Fig. 4B). Peaks began to shift when stochasticity was strong enough to move population vectors occasionally at least as far from c_1 and c_2 as *e*. Linearization theory failed when this occurred because linearization at points in the attractor does not capture the dynamical behavior of the LPA model at points farther from the attractor than *e*.

For weak stochasticity, the SD-LPA model initialized with a population vector near c_i will produce in one time step a vector nearer c_j ($j \neq i$) than c_i. Stronger stochasticity will cause the model to occasionally produce a vector closer to c_i, shifting the phase of the n.f. 1 oscillation of the system (Henson *et al.* 1998). The relative frequency of phase shifting, as a function of the Σ factor, becomes appreciably greater than 0 for the same range of Σ factors at which stochasticity becomes strong enough to move populations further from the attractor than *e*, and peak motion begins (Fig. 11C).

Figure 11: Explanation of peak motion patterns with increasing stochasticity for (A) $c_{pa} = 0.5$ and (B, C) $c_{pa} = 0$ in the SD-LPA model. First panel of (A): for every square-root-scale experimental population vector $\mathbf{x}_t$ with both $\mathbf{x}_t$ and $\mathbf{x}_{t+3}$ within 10 units of $\mathbf{a}_1$ (text), the components $\mathbf{h}_t$ and $\mathbf{h}_{t+3}$ of $\mathbf{x}_t - \mathbf{a}_1$ and $\mathbf{x}_{t+3} - \mathbf{a}_1$ in the direction of $\mathbf{v}_1$ were computed, and the signed length (s.l.) of $\mathbf{h}_{t+3}$ was plotted against the signed length of $\mathbf{h}_t$. Other panels of (A): same as the first panel but using $\mathbf{a}_2$, $\mathbf{v}_2$ and $\mathbf{a}_3$, $\mathbf{v}_3$ respectively. Plotted linear regression slopes were significantly negative (99% level), and took values greater than -1, demonstrating over-compensatory decay of perturbations from $\mathbf{a}_\mathrm{i}$ in the direction of $\mathbf{v}_\mathrm{i}$ in three time steps and supporting the mechanism by which theory explains the peak at n.f. 0.33 for $c_{pa} = 0.5$. (B) The increasing trend gives the mean, across length-1024 SD-LPA-model-generated time series, of the distance between each population vector and the attractor comprised of c_1 and c_2. Vertical bars stretch between the 2.5^{th} and 97.5^{th} percentiles; they extend beyond the axes for large Σ factors. The solid horizontal line gives the distance between *e* and the attractor. Means of distances between experimental population triples and the attractor were 47.7, 51.2, and 48.1 for the three replicates with $c_{pa} = 0$, agreeing with simulated values. (C) We used the SD-LPA model with each of a range of Σ factors to generate 50 time series of length 1024. For each time series, the percentage of steps at which populations failed to switch from being closer to c_1 to being closer to c_2 or *vice versa* was calculated. The increasing trend gives mean percentages across the 50 time series; vertical bars span the 2.5^{th} and 97.5^{th} percentiles. *See separate file for the figure.*

Supplementary References:

1. Brillinger, D. (2001) *Time Series: Data Analysis and Theory*. Society for Industrial and Applied Mathematics, Philadelphia.
2. Cushing, J.M., Costantino, R.E., Dennis, B., Desharnais, R.A. & Henson, S.M. (2003) *Chaos in Ecology: Experimental Non-linear Dynamics*. Academic Press, New York.
3. Dennis, B., Desharnais, R.A., Cushing, J.M., Henson, S.M., Costantino, R.F. (2001) Estimating chaos and complex dynamics in an insect population. *Ecological Monographs* **71**(2), 277-303.
4. Greenman, J. V. & Benton, T. G. (2005) The frequency spectrum of structured discrete-time population models: Its properties and their ecological implications. *Oikos*, 110, 369-389.
5. Henson, S. M., Cushing, J. M., Costantino, R. F., Dennis, B., & Desharnais, R. A. (1998) *Proceedings of the Royal Society of London* B, 265, 2229-2234.
6. King, A.A., Costantino, R.F., Cushing, J.M., Henson, S.M., Desharnais, R.A., Dennis, B. (2004) Anatomy of chaotic attractor: subtle model-predicted patterns

revealed in population data. *Proceedings of the National Academy of Sciences of the U.S.A.* **101**, 408-413.

7. Wiesenfeld, K. (1985) Noisy precursors of nonlinear instabilities. *Journal of Statistical Physics*, 38, 1071-1097.