Research

Oikos 116: 1726–1736, 2007
doi: 10.1111/j.2007.0030-1299.16101.x,

Subject Editor: Esa Ranta, Accepted 1 June 2007

Environmental colour intensifies the Moran effect when population dynamics are spatially heterogeneous

David A. Vasseur

D. A. Vasseur (dvasseur@ucalgary.ca), Dept of Biological Sciences, Univ. of Calgary, Calgary, AB, Canada, T2N 1N4.

Evidence for synchronous fluctuations of spatially separated populations is ubiquitous in the literature, including accounts within and across taxa. Among the few mechanisms explaining this phenomenon is the Moran effect, whereby independent populations are synchronized by spatially correlated environmental disturbances. The body of research on the Moran effect predominantly assumes that environmental disturbances within a local site are serially uncorrelated; that is, successive observations in time at a particular local site are independent. Yet, many environmental variables are known to possess strong temporal autocorrelation – a character which has often been described as 'colour'. The omission of environmental colour from research on the Moran effect may be due in part to the lack of methods capable of generating sets of time series with a desired colour and spatial correlation. Here I present a novel and simple method designated as 'phase partnering' to generate such sets of time series and I investigate the combined impact of spatial correlation and environmental colour on population synchrony in two common models of population dynamics. For linear population dynamics, and for a subset of nonlinear population dynamics, coloured environments intensify the Moran effect when population dynamics are spatially heterogeneous; in coloured environments the spatial correlation between populations more closely mimics the spatial correlation between their respective environments. Given that most environmental variables are coloured, these results imply that the Moran effect may be a far more significant driver of regional-scale population and interspecific synchrony than is currently believed.

Spatial population synchrony is a well-known and seemingly ubiquitous phenomenon, occurring on local and regional scales within a wide range of different taxa (reviewed by Bjørnstad et al. 1999, Liebhold et al. 2004, Ranta et al. 2006). Three mechanisms are known to give rise to this phenomenon: 1) extrinsic fluctuations in climate variables which are spatially-correlated (Moran 1953, Royama 1992, Hudson and Cattadori 1999); 2) dispersal of individuals between spatially segregated populations (Heino 1998, Lande et al. 1999, Ranta et al. 1999, Ripa 2000); and 3) interactions with other populations which are synchronized across space as a consequence of 1), 2), or due to a homogenous and large geographic range (Small et al. 1993, Ims and Andreassen 2000, Koenig 2001, Liebhold et al. 2004, Ranta et al. 2006, Ripa and Ranta 2007). A special case of the first mechanism was formally described by Moran (1953) and has since become know as Moran's theorem (Royama 1992). Moran's theorem states that a pair of spatially separated populations, obeying identical linear dynamics, in absence of dispersal, and subject to random climate fluctuations, will demonstrate a cross-correlation that is equivalent to the spatial correlation of their climates. The lack of such ideal conditions in the field has given rise to a more loosely applied term – the 'Moran effect' – which implies a climate induced synchronization that is weaker than predicted by Moran's theorem (Royama 1992). Given that climatic fluctuations are spatially autocorrelated over large distances (Koenig 2002), the Moran effect has been suggested as a significant driver of population synchrony (Hudson and Cattadori 1999).

Since Royama (1992) revisited Moran's (1953) early work, theoretical research on the Moran effect has reawakened with particular attention given to testing the constraints of Moran's original theorem. In particular researchers have evaluated the Moran effect in populations with nonlinear dynamics (Ranta et al.

1997, 1999, Grenfell et al. 1998, Greenman and Benton 2001, Engen and Sæther 2005, Royama 2005, Abbott 2007), dispersal (Ranta et al. 1995, 1999, Kendall et al. 2000, Ripa 2000, Liebhold et al. 2006, Abbott 2007), and in populations whose dynamics are governed by different density dependences (Peltonen et al. 2002, Ripa and Ives 2003, Engen and Sæther 2005, Royama 2005, Hugueny 2006, Liebhold et al. 2006). This latter constraint of Moran's theorem may be particularly important given that geographic variation may lead to qualitatively different population dynamics within a species (Peltonen et al. 2002, Hugueny 2006, Liebhold et al. 2006) and that interspecific synchrony occurs (Liebhold et al. 2004, Tedesco et al. 2004, Ranta et al. 2006) despite obvious and often large differences in population dynamics. Recent work has shown that the Moran effect generally weakens as the parameters governing density dependence diverge (Ripa and Ives 2003, Engen and Sæther 2005, Royama 2005, Hugueny 2006), suggesting that it may be a limited source of synchrony when dynamics vary distinctively across space.

Previous work on population synchrony has almost exclusively assumed that climate fluctuations within a local habitat lack temporal autocorrelation (Heino 1998, Fontaine and Gonzalez 2005), yet many climate variables are known to demonstrate positive temporal autocorrelation – a positive relationship between successive observations recorded at a fixed location (Halley 1996, Pelletier 2002, Cyr and Cyr 2003, Vasseur and Yodzis 2004). This character can be described by the spectral exponent (γ), which expresses the relative contribution of fluctuations at different frequencies to the temporal variance of the observations, assuming that these contributions scale as $1/f^{\gamma}$ (Halley 1996). When $\gamma = 0$, all frequencies contribute equally to the temporal variance and successive observations are unrelated (they lack autocorrelation); this form of disturbance is commonly referred to as 'white' noise. Alternatively, when $\gamma > 0$, subsequent observations are positively associated due to an increased importance of low-frequency contributions to the temporal variance. This family of disturbances are commonly referred to as 'red' or 'reddened' noises in analogy to the frequency composition of red light (cf. 'blue' noises which have $\gamma < 0$) and are well suited to describing temporal fluctuations in environmental variables (Halley 1996).

Although the importance of coloured noise for population and community dynamics has been well studied (Ripa and Lundberg 1996, Kaitala et al. 1997, Petchey et al. 1997, Heino 1998, Cuddington and Yodzis 1999, Ripa and Ives 2003, Wichmann et al. 2005, Schwager 2006), few studies have described its impact on population synchrony. Heino (1998) demonstrated that red noise could improve population synchrony in a non-linear metapopulation model where patches were linked by dispersal and patches experienced a 'local' and 'global' set of 'reddened' environmental conditions. Although this formulation introduced spatial correlation among the coloured environments in each local habitat, it did not allow the correlation to vary with the distance between patches as research suggests it should (Koenig 2002). Using autoregressive models, Ripa and Ives (2003) investigated the impact of coloured noise on the dynamics of competitors and predator–prey systems, but limited their discussion to the analytically solvable cases where the environmental noises were either independent or identical. They showed that environmental colour caused dynamics to align along certain eigenvectors in the phase space, therefore altering the correlation between populations. A recent laboratory experiment demonstrated that environmental fluctuations possessing a reddened spectrum generated synchronous dynamics in rotifer populations whereas environmental fluctuations possessing a white spectrum did not (Fontaine and Gonzalez 2005). Experimental work has outpaced theory in this area, potentially, for want of a method to generate sets of coloured noises with any desired cross- (spatial) correlation. Herein I introduce a novel and simple method designated as 'phase partnering' for generating such noises and I test two commonly employed population models (linear and non-linear) to determine the influence of environmental colour on the operation of the Moran effect.

Generating coloured cross-correlated noise using 'phase partnering'

The goal of the method outlined below is to generate a pair of 'environmental' time series which have a desired cross-correlation (herein referred to as ρ_{ξ}) and a desired colour (referred to precisely by the spectral exponent γ). Although multiple methods exist for independently generating coloured and cross-correlated time series (Halley 2004, Ripley 2006), sequential application of such methods demonstrates the inherent problem; post-hoc colouring of cross-correlated noises (using spectral synthesis or autoregressive filtering) introduces large variability into the desired cross-correlation (below), and vice-versa, post-hoc cross-correlating of coloured noises (e.g. using linear combination) introduces large variability into the desired colour (Vasseur unpubl.). Below I outline an approach which addresses this problem by concurrently applying the desired colour γ and cross-correlation ρ_{ξ} to a pair of environmental time series. With this method, the desired colour and cross-correlation of the time series are achieved with near-perfect accuracy and precision.

Environmental time-series $\xi_{1,t}$ and $\xi_{2,t}$ are generated using a modification of the method known as spectral synthesis (Cuddington and Yodzis 1999, Halley 2004, Fontaine and Gonzalez 2005) in which every tth observation is generated by summing n/2 sinusoids:

$$\xi_{i,t} = \sum_{f=1}^{n/2} \frac{1}{f^{\gamma/2}} \sin[2\pi ft/n - 2\pi ft/n + \theta_i(f)] \quad (1)$$

where n is the number of observations in the time series, t is time, f is frequency, and $\theta_i(f)$ is a vector of phase operators which are independent and randomly drawn from a uniform distribution on $[0,2\pi)$. This procedure generates stationary time series (for $\gamma < 2$) which are normally distributed with zero-mean and a spectral exponent γ. Multiple iterations of this method produce time series which are independent – due to the independence of phase vectors θ in Eq. 1. However, if the phase vectors of successive iterations of Eq. 1 are not independent, then neither are the resultant time series ξ_1 and ξ_2 – an aspect which the method outlined below utilizes to generate cross-correlation.

The cross-correlation between two time series ξ_1 and ξ_2 can be predetermined by 'partnering' the phase vectors θ_1 and θ_2 so that the cross-correlation between the two time series is preserved at each of the n/2 frequencies. To generate any desired cross-correlation between ξ_1 and ξ_2, the phase operators $\theta_i(f)$ are related by the function:

$$\theta_2(f) = \theta_1(f) + \delta \quad (2)$$

where δ is constant phase shift which is determined by ρ_ξ according to $\delta = \cos^{-1}(\rho_\xi)$ (see Appendix 1 for derivation). Under this transformation, when $\rho_\xi = 1$, the phase shift $\delta = 0$ and thus ξ_1 and ξ_2 are identical time series. At the alternative extreme $\rho_\xi = -1$, the phase shift $\delta = \pi$ and ξ_1 and ξ_2 are simply reflections of each other about the origin. For the intermediate cases, where $-1.0 < \rho_\xi < 1.0$, ξ_1 and ξ_2 are unique, however, the addition of a constant phase shift δ to each sinusoidal component generates a tendency for the first time series to lag behind the second. This tendency is undesirable since it can introduce large time-lagged cross-correlations between environmental and population dynamics. To ensure that ξ_1 does not consistently lag behind ξ_2 it is possible to make use of the fact that $\cos^{-1}(\rho_\xi)$ is a multivalued function and randomize the phase shift δ according to $\delta(f) = \varepsilon_f \times \cos^{-1}(\rho_\xi)$ where ε_f is randomly selected from a 'coin toss' as -1 or 1 with equal probability. This randomization preserves the cross correlation between ξ_1 and ξ_2 but shuffles the 'leading' and 'lagging' effect amongst frequencies so that there is no apparent effect between the two time-series. In practice, spectral exponents with large magnitude may be prone to this behaviour even when the lags are shuffled amongst frequencies, since a small proportion of frequencies and their associated lags, dominate the variance of the resultant time series. For the range of spectral exponents used herein, and those most typical of natural environments, this 'small sample' issue is not problematic. Following generation of the time-series, their variances can easily be standardized to a desired value by vector multiplication or by the method of spectral mimicry (Cohen et al. 1999) which ensures that the mean, variance, and range of ξ_1 and ξ_2 are equal, but which can add some variability to the spectral exponent and cross-correlation of the noises (Vasseur unpubl.).

To ensure that the method itself does not introduce any artifactual elements into the time series it is reasonable to analyze the two series in the absence of any cross-correlation and colour (ρ_ξ, $\gamma = 0$; note that in practice it is unreasonable to generate independent vectors of white noise using phase partnering when a simpler random number generator would suffice). In this example the sample mean environmental conditions $\hat{\mu}_{\xi_1}$ and $\hat{\mu}_{\xi_2}$ equal zero and the variances are standardized to a value of 1 (Fig. 1). The frequency distributions of each series do not differ from normally distributed noise (Kolmogorov-Smirnov test; $D = 0.023$, 0.028; $p > 0.4$) and the autocorrelation function (ACF; Fig. 1d) shows no apparent autocorrelation within either series (Durbin-Watson test; $d = 2.009$, 1.999). The cross-correlation function (CCF; Fig. 1c) demonstrates that fewer than 5% of the cross-correlation coefficients deviate significantly from zero, and more importantly, that none lie far outside the significance intervals $\pm 2/\sqrt{n}$ (Chatfield 2004), indicating that there is no appreciable correlation at any time-lags between the two time series.

In Fig. 2 the same plots are shown for two series with $n = 1000$, $\gamma = 1.0$, and $\rho_\xi = 0.75$ Comparing Fig. 1a and 2a, the temporal autocorrelation that is introduced by the change in γ is evident in the relatively small changes at successive time steps. Like the previous example the frequency distributions do not differ from normal ($D = 0.016$, 0.019; $p > 0.88$), however the test for autocorrelation obviously fails due to the autocorrelation introduced into each series by the non-zero spectral exponent γ (Fig. 2d). Here the CCF (Fig. 2c) also demonstrates a long tailed distribution due to the positive cross-correlation between- and autocorrelations within these series (Fig. 2d). The most important observation from Fig. 2d is that the cross-correlation is maximized at lag–0, where the desired cross-correlation coefficient is near-perfectly conserved ($\hat{\rho}_\xi = 0.7501$).

Like 1/f noises, first order autoregressive models AR(1) are also commonly used to generate autocorrelated noises for the subsequent study of population dynamics in noisy environments (Petchey et al. 1997, Heino 1998, Ripa and Ives 2003, Wichmann et al.

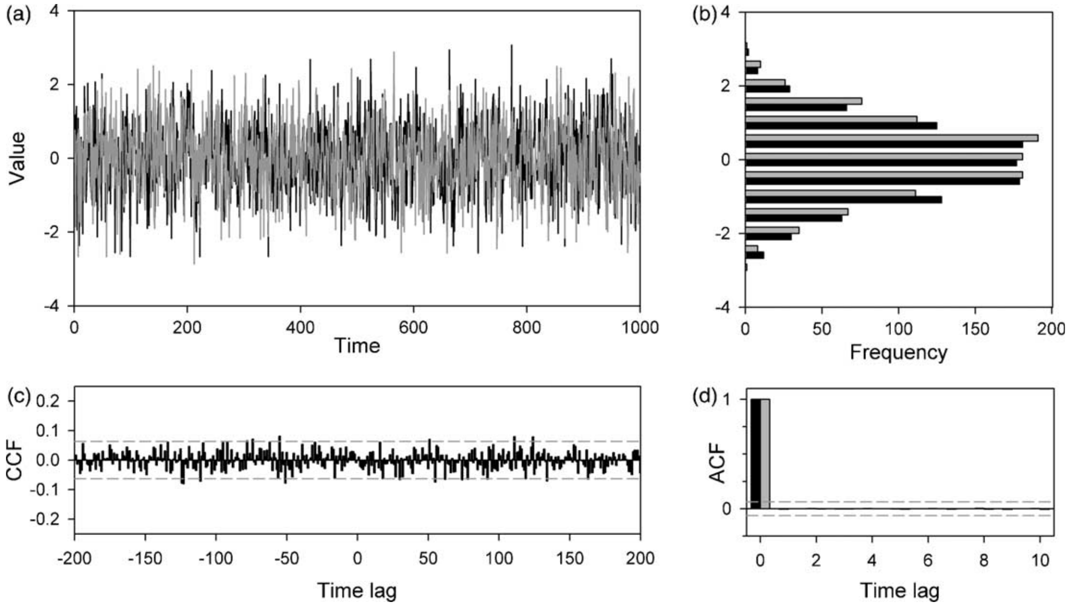

Fig. 1. Time series analysis of ξ_1 (black lines) and ξ_2 (gray lines) for $\gamma = 0$, $\rho_\xi = 0$, and $n = 1000$, showing (a) the time series of the two noise processes; (b) frequency distributions of the values from the two processes; (c) the cross-correlation function (CCF); and (d) the autocorrelation functions (ACF) of the two processes. Values in excess of the dashed lines ($\pm 2/\sqrt{n}$) in panels (c) and (d) are significantly different from zero ($\alpha = 0.05$).

2005, Schwager et al. 2006). AR(1) noise differs in many respects to 1/f noise (Halley 1996); most notably, AR(1) noise has a finite variance (Wichmann et al. 2005) whereas the variance of 1/f noise increases with the length of the series (Halley 1996). Despite clear differences in the two models, there is still some

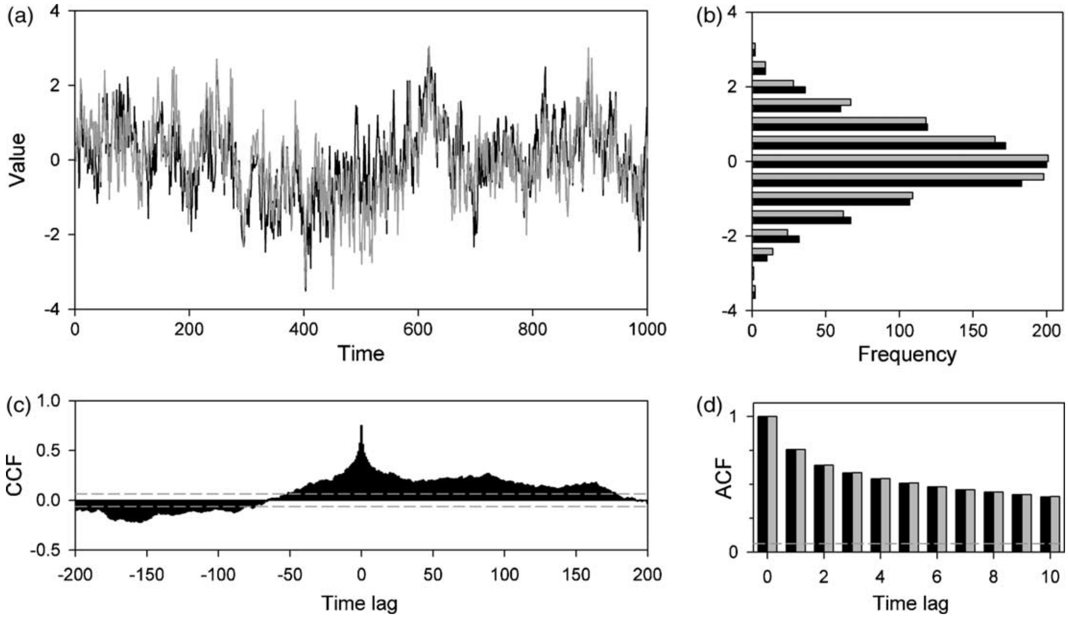

Fig. 2. Time series analysis of ξ_1 (black lines) and ξ_2 (gray lines) for $\gamma = 1.0$, $\rho_\xi = 0.75$, and $n = 1000$, showing (a) the time series of the two noise processes; (b) frequency distributions of the values from the two processes; (c) the cross-correlation function (CCF); and (d) the autocorrelation functions (ACF) of the two processes. Values in excess of the dashed lines ($\pm 2/\sqrt{n}$) in panels (c) and (d) are significantly different from zero ($\alpha = 0.05$).

debate over which model best describes empirical records of environmental variability (Cyr and Cyr 2003, Vasseur and Yodzis 2004); however, inferential climate records suggest that power-law (1/f) relationships prevail over the long-term (Koscielny-Bunde et al. 1998, Pelletier 2002). AR(1) noises can also be constructed to produce a desired level of cross-correlation (ρ_φ) and so comparing cross-correlated AR(1) noises with those produced by 'phase partnering' provides a useful contrast. AR(1) environmental time series is generated by:

$$\varphi_i(t) = \alpha \cdot \varphi_i(t-1) + \varepsilon_i(t) \qquad (3)$$

where $0 \leq \alpha \leq 1$ controls the degree of autocorrelation, $\varepsilon_i(t)$ is random variable drawn from a normal distribution, and i is an index variable. If two time-series are generated using two random vectors (ε_1 and ε_2) which are correlated (for such methods see Ripley 2006), the resultant time series have an expected cross correlation ρ_φ which is equal to that of the random vectors. However, filtering these correlated noises through the AR(1) process can introduce relatively large errors into the resultant cross correlation. Figure 3 compares the error introduced by this method compared to phase partnering as the autocorrelation (colour) increases. For AR(1) noise the error grows as the as the autocorrelation increases; this occurs because the two series become increasingly dominated by historical (and perhaps different) values. The origin of the pattern shown for phase-partnered 1/f noise (Fig. 3b) is less clear, but may be a consequence of fewer frequencies contributing to the noise at larger values of γ. Both models produce the largest errors when the desired cross-correlation is zero (not shown). When comparing the two models in uncorrelated environments (α, $\gamma = 0$) the error is approximately an order of magnitude smaller for phase-partnered 1/f noise; however, for in strongly autocorrelated noises ($\alpha = 1$, $\gamma = 2$) the error is four to five orders of magnitude smaller.

Population dynamics in coloured cross-correlated environments

Although the method described above is capable of generating time-series over the entire range of noise colours, the remainder of this paper focuses on the impact of reddened environmental noise ($\gamma > 0$), since environmental variables of ecological interest almost exclusively take such form (Pelletier 2002, Vasseur and Yodzis 2004, Cyr and Cyr 2003). Herein 'colour' refers to reddened noises.

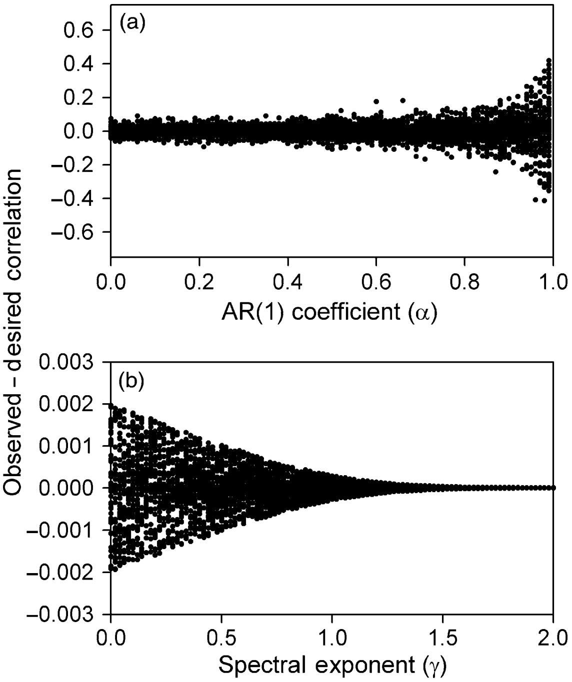

Fig. 3. The error (observed − expected correlation) produced by (a) an AR(1) process and (b) phase association over a range of autocorrelation (colour). At each value of the independent variables (α and γ), a uniform set of cross-correlations from −1.0 to 1.0 were used to generate the points on the figure.

Linear population dynamics

Consider a pair of biological populations whose log-transformed abundances are defined Y_1 and Y_2 and whose dynamics are governed by a first-order linear autoregressive equation:

$$Y_{i,t+1} = a_i + b_i\, Y_{i,t} + \xi_{i,t} \text{ where } i = 1, 2 \qquad (4)$$

The parameters a and b describe the population renewal process and ξ_1 and ξ_2 are environmental variables. Here, ξ_1 and ξ_2 are generated by the method of phase partnering described above and standardized to have zero mean, equal variances, a colour (γ) and a cross-correlation at lag-0 represented by ρ_ξ. Using an analogous second-order model, Moran (1953) noted that when Y_1 and Y_2 have identical density dependences, their cross-correlation is identical to that of their environments. For Eq. 4, Moran's theorem states that when $b_1 = b_2$, the proportionality constant $K = \rho_Y/\rho_\xi$ will equal 1. Since it is unlikely that two spatially separated populations would have identical density dependences, Royama (2005) and Hugueny (2006) expanded Moran's original theorem to populations with spatially heterogeneous dynamics. In this case, K is always less than 1 and as populations become increasingly heterogeneous, their synchrony weakens (K decreases; Royama 2005). Note that Hugueny (2006) describes a similar proportionality constant which is termed the demographic component (DC) of population synchrony.

The colour of environmental noise has no influence on populations with homogeneous dynamics (i.e. identical density dependence); however, as the population dynamics become increasingly heterogeneous (as $|b_1 - b_2|$ grows), noise colour becomes an important determinant of population synchrony (Fig. 4). For most parameter combinations, coloured environments ($\gamma > 0$) maintain higher levels of population synchrony than do white environments ($\gamma = 0$). As the environmental noise colour becomes deeply reddened ($\gamma \approx 2$) the population cross-correlation nearly equals that of the environment, demonstrating that the Moran effect can synchronize populations whose dynamics are governed by very different density dependent processes.

Coloured environments are characterized by long runs of similar conditions due to the relatively weak influence of high frequencies in their spectrum. This property allows populations governed by models with stable equilibria to maintain a closer association with their equilibrium densities – a phenomenon known as ‘tracking’ (Roughgarden 1975, May 1976, Petchey 2000). For any pair of populations whose (stable) equilibria demonstrate a correlated response to changing environmental conditions, improved ‘tracking’ should thereby cause population fluctuations to mimic the correlated fluctuations the environment. In this model (Eq. 4) the equilibrium varies with environmental conditions according to $Y_t^* = (a + \xi_t)/(1 - b)$, and it is stable provided that population dynamics are stationary $-1 < b < 1$ (Royama 2005, Hugueny 2006). If the mean environmental condition is zero, then the average density around which the equilibrium fluctuates is $a/(1 - b)$ and the residual fluctuations are equal to $\xi_1/(1 - b)$. Provided that populations are subject to the stationarity constraint, the residuals, and thereby the equilibrium densities will have the same cross-correlation as that of their environments (note that the mean environmental condition need not be zero for this to be true). Thus, the improved tracking ability afforded by coloured environments intensifies the Moran effect because the environmental correlation is directly transferred to fluctuations in population equilibria. Although this result holds for any pair of populations subject to the stationary constraint, the quantitative increase in synchrony can vary largely in parameter space; as b nears the boundary conditions for stationarity the return tendency of the population becomes very weak. This severely limits the population's ability to track environmental conditions and its potential to correlate with other non-identical populations (Fig. 4; Ripa and Ives 2003, Hugueny 2006). Thus, when one population in a pair has a very weak return tendency, coloured noise has a reduced impact on the Moran effect.

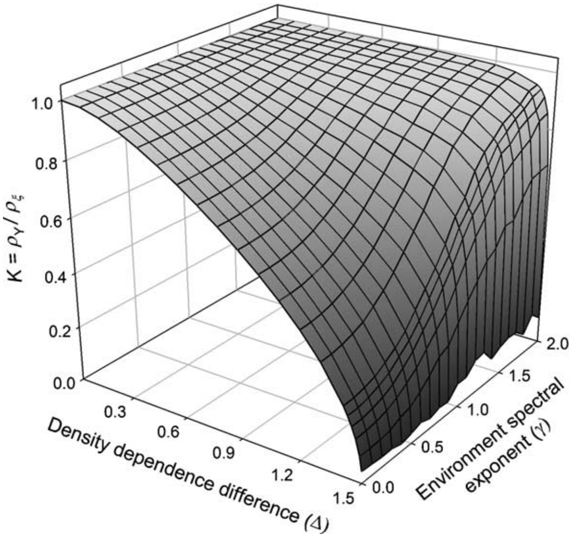

Fig. 4. The impact of coloured noise on the synchrony of linear population dynamics. The proportionality constant $K = \rho_Y/\rho_\xi$ is shown for a range of difference in density dependence (Δ) and environmental spectral exponents (γ). Population dynamics are governed by a first-order autoregressive model (Eq. 3) where $b_1 = -0.5$, $b_2 = b_1 + \Delta$, and $a_{1,2} = 0$ (although the choice of a has no bearing on the results). When $\Delta = 1.5$ the dynamics of Y_2 are non-stationary. In the white noise environment ($\gamma = 0$) the response of K to the difference in density dependence reduces to the closed-form solution $K = [(1 - b_1)^2(1 - b_1)^2]^{1/2}(1 - b_1 b_2)^{-1}$ (modified from Royama 2005, Hugueny 2006). K was calculated at each intersection of the grid lines from 100 independent replicate simulations of Eq. 4 with $\mu_\xi = 0$ and $\sigma_\xi = 0.1$. Initial Y values were chosen randomly between 0 and 1. The model was iterated for 1000 time-steps but only the latter 500 time-steps were used to estimate ρ_Y.

Non-linear population dynamics

A similar but non-linear analogue to the above model (Eq. 4) is the discrete Ricker model. For two populations, $i = 1, 2$, and including environmental noise, population abundance is calculated as (Greenman and Benton 2001)

$$Y_{i,t+1} = Y_{i,t} \times e^{(a_i - Y_{i,t} + \xi_{i,t})} \tag{5}$$

where a_i is the intrinsic rate of increase. The dynamic behaviour of this model varies with the parameter a, transitioning from a monotonic approach to equilibrium ($a < 1$), to damped oscillations ($a < 2$), to unstable cycles ($a < 2.69$), and eventually to chaos in the absence of any environmental noise. It is rather obvious from the model's formulation that changes in the environmental condition $\xi_{i,t}$ have an effect that is equivalent to changing a, thus, environmental fluctuations can induce short-term changes in the dynamic behaviour of the model. Increasing the standard deviation of environmental noise σ_ξ increases the potential for the equilibrium to cross these bifurcation thresholds as a consequence of environmental conditions. Since two populations whose environments are correlated (but not identical) may, at any time t, be subject to different underlying dynamics (e.g. cycles vs chaos), the spatial correlation between populations tends to be well below that of their environments – even for populations with identical density dependences (Greenman and Benton 2001). Increasing the standard deviation of environmental noise exacerbates this difference (Greenman and Benton 2001) and the lack of a consistent ratio between population and environmental correlation precludes use of the proportionality constant K to depict the results (as in the linear model).

As in the linear model, coloured environments improve the synchrony between Y_1 and Y_2 when dynamics are spatially heterogeneous (Fig. 5), albeit for only a subset of parameter space and subject to the underlying dynamics of the model. In Fig. 5a and 5b, the dynamics of one population are fixed in the stable parameter regime while those of the second population

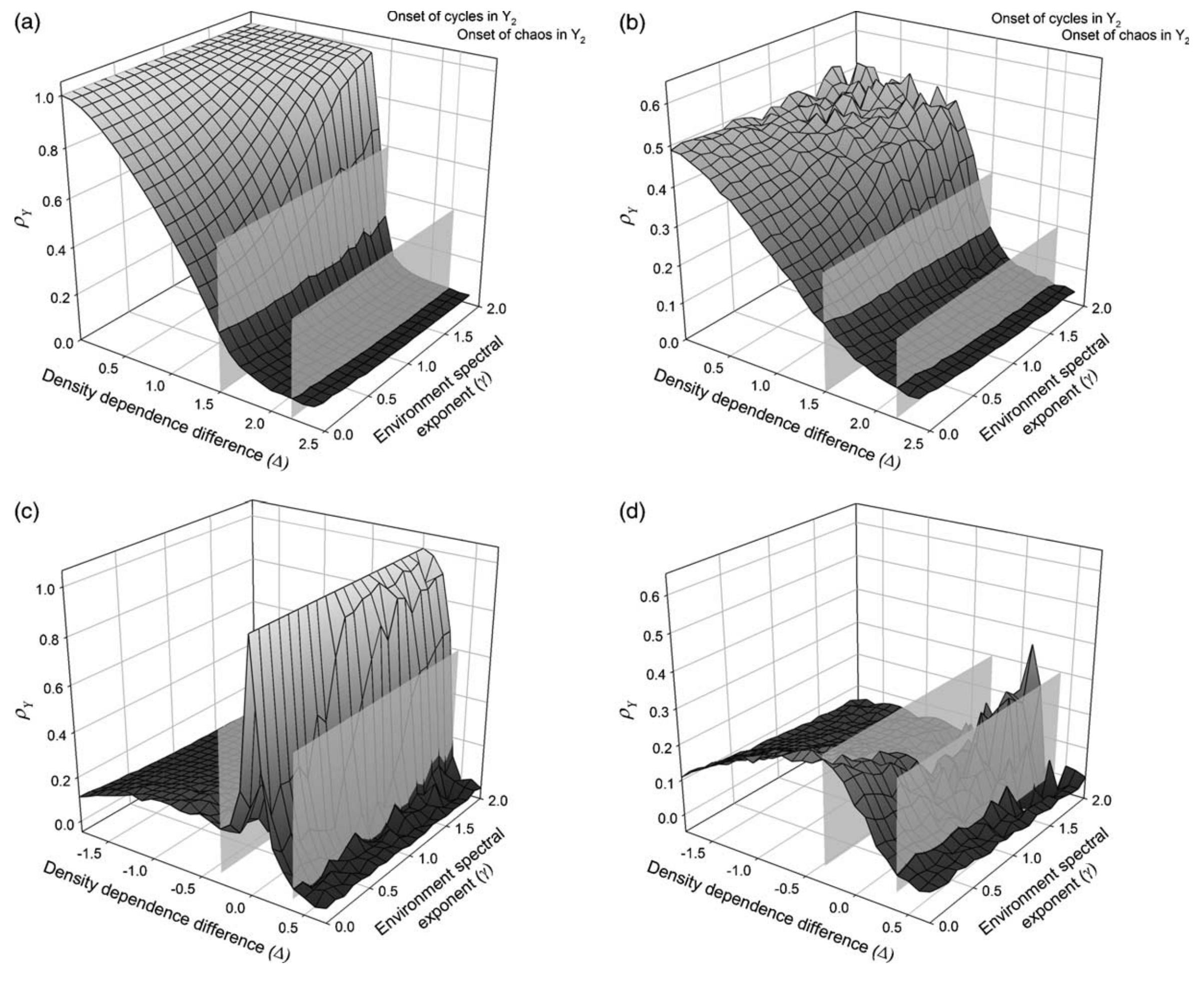

Fig. 5. The impact of coloured noise on the synchrony of non-linear population dynamics. The cross correlation of population dynamics ρ_Y is shown for a range of difference in density dependence (Δ) and environmental spectral exponents (γ) given $\rho_\xi = 1.0$, $\sigma_\xi = 0.1$ (a) and (c) $\rho_\xi = 0.5$, $\sigma_\xi = 0.2$ (b) and (d). Population dynamics are governed by a Ricker model (Eq. 4) where $a_1 = 0.5$, and $a_2 = a_1 + \Delta$. The intersecting gray planes correspond to values of Δ at which the dynamics of Y_2 cross bifurcation thresholds. ρ_Y was estimated at each intersection of the grid lines from 100 independent replicate simulations of Eq. 5. Initial Y values were chosen randomly between 0 and 1 [but in panels (c) and (d), $Y_1(0) = Y_2(0)$]. The model was iterated for 1000 time-steps but only the latter 500 time-steps were used to estimate ρ_Y.

are varied (across stable and unstable parameter space). When both populations are in the stable parameter regime (Fig. 5a-b) the results mimic those of the linear model (Fig. 4); for spatially heterogeneous population dynamics, coloured environmental noise intensifies population synchrony. However, this outcome is limited to the fraction of parameter space where both populations are well bounded from unstable dynamics. The synchronizing potential of coloured noise breaks down quickly once the density dependence difference (Δ) reaches a critical threshold approximately $2\sigma_\xi$ below the bifurcation from damped oscillations to unstable cycles (Fig. 5; cf. panels a and b which have a different σ_ξ). Beyond this threshold, improved environmental tracking no longer benefits population synchrony as the underlying dynamic of Y_2 is, with increasing frequency, governed by a qualitatively different equilibrium behaviour. When the environmental correlation is reduced and the distribution of fluctuations is increased (Fig. 5b), the population correlation becomes increasingly variable, however, the effect of environmental colour is still visible and its mode of influence the same.

For the non-linear model (Eq. 5), it is also noteworthy to examine the impact of coloured noise when one population is fixed in the unstable (cyclic) parameter regime while the second is varied (Fig. 5c-d). As in Fig. 5a and Fig. 5b, when the two populations are governed by different equilibrium behaviours, environmental colour offers no improvement in population synchrony. When the intrinsic dynamics of both populations are cyclic, there is an additional barrier to synchrony; frequent switches occur between states where fluctuations are exactly in-phase (synchronized) or exactly out-of-phase (Greenman and Benton 2001). Even when environments are identical, certain initial conditions can generate perfectly out-of-phase population dynamics. To control for this, the initial densities of the two populations are made identical in Fig. 5c and 5d, ensuring that populations remain in the 'in-phase' state when environments are perfectly correlated. Under this constraint, environmental colour improves population synchrony over a small range of heterogeneous dynamics (Fig. 5c); however, when the environmental correlation is moderate ($\rho_\xi = 0.5$) and the size of the noise distribution is increased, mixing of the in- and out-of-phase states supersedes any impact that colour has on population synchrony (Fig. 5d).

Discussion

In both linear and non-linear models, coloured environments can generate a substantial increase in population synchrony over white environments when population dynamics are spatially heterogeneous. For the linear population model investigated, synchrony increases in coloured environments when dynamics are spatially heterogeneous (when populations are identical Moran's theorem is always upheld). Strongly reddened environmental noise (i.e. $\gamma \approx 2$) overrides nearly all differences in density dependences and maintains the population correlation very near to the environmental correlation. Only when one population has a very weak return tendency does coloured noise have little impact upon population synchrony. For the nonlinear (Ricker) model explored, population synchrony increases in coloured environments when the underlying dynamics of the two populations are both stable. For a pair of cyclic populations, coloured noise can improve synchrony provided that the environmental correlation is high and initial population densities favour in-phase oscillations. Coloured noise offers no benefit to population synchrony when the two populations are governed by different qualitative dynamics. Given the increase in population synchrony offered by coloured noise for linear – and a subset of nonlinear dynamics, the Moran effect may be a far more significant driver of synchrony than is currently believed.

Theory predicts that 'tracking', the association between population fluctuations and their stochastic equilibria, should be improved in reddened environments because environmental fluctuations are shifted to scales at which populations can better respond (Roughgarden 1975, May 1976). This phenomenon has been confirmed in laboratory microcosms (Petchey 2000, Laakso et al. 2003), albeit by demonstrating an increased correlation between population and environmental (rather than equilibrium) fluctuations in coloured environments. While one must assume a monotonic relationship between environmental and equilibrium fluctuations in order to accept these conclusions, the influence of tracking on population synchrony is far less restrictive; improved tracking will lead to increased synchrony between populations provided only that their equilibria respond coherently (not necessarily monotonically) to environmental fluctuations. Here, experiments have outpaced theory. Rotifer populations (of the same species) demonstrated synchronous dynamics in experimental microcosms when resource fluctuations followed a red noise regime and lacked synchrony when resource fluctuations were white (Fontaine and Gonzalez 2005). This result occurred despite an identical environment across treatments, demonstrating that populations had little success tracking fluctuations in white environments.

Population synchrony is most often explained as a consequence of the Moran effect and dispersal acting in concert (Lande et al. 1999, Ranta et al. 1999, Kendall et al. 2000, Ripa 2000). However, the combined effects of dispersal and heterogeneity in density dependence between patches have not yet been investigated

(Hugueny 2006). Understanding how coloured environmental noise impacts population synchrony in this setting is an avenue warranting further research. In rare instances, the role of dispersal for synchrony can be neglected and the potential importance of coloured noise for synchrony can be surmised. Grenfell et al. (1998) described the dynamics of feral sheep on two nearby islands in the St. Kilda archipelago using a nonlinear model. They found that the observed level of population synchrony ($r = 0.685$) required what they considered to be an unsupported degree of environmental correlation (greater than $r = 0.9$). Although Koenig (2002) later pointed out that this degree of environmental correlation was indeed plausible, given evidence of similar high correlations at nearby sites, it stands to reason that environmental colour could also account for the observed level of synchrony while requiring a lesser degree of environmental correlation.

Ecologists have recently recognized the need to reject white noise as a model environment because many environmental variables, including temperature, precipitation, humidity, river height, and some seasonal indices, have significantly reddened spectra (Koscielny-Bunde et al. 1998, Cuddington and Yodzis 1999, Pelletier 2002, Cyr and Cyr 2003, Vasseur and Yodzis 2004). Although there appears to be variation in the extent of reddening across different time-scales, red noise pervades environmental time series at temporal scales ranging from a few days to hundreds of decades (Pelletier 2002; cf. Koscielny-Bunde et al. 1998). This suggests that noise colour should be equally important for the synchrony of species whose dynamics are influenced by short-term environmental conditions (e.g. daily temperature or irradiance), and those which are determined by long-term environmental conditions (e.g. winter severity or annual rainfall). The obvious and rather difficult task of linking the colour of environmental noise to synchrony in natural populations lies in determining exactly how population fluctuations are linked to environmental fluctuations; understanding which environmental variables are important for a particular population at each stage of its life cycle will ultimately lead to a better understanding of how synchrony is achieved.

The method of 'phase partnering' described herein provides a robust and simple framework for generating coloured cross-correlated noises and it can be easily extended to generate larger sets of noise vectors. This property may be particularly useful for research into metapopulation synchrony, synchronous traveling waves (Ranta et al. 1999, 2006), and for food web dynamics in varying environments – given that different species may be sensitive to different, but correlated, environmental variables. In the two patch system described in this study, the method requires only the specification of a single cross-correlation parameter. However, in an n-patch system the number of cross-correlations grows as $n(n-1)/2$ and not all of these parameters can be freely specified, rather they are subject to the general constraints of correlation matrices. For example, in a set of three patches (x, y, and z) each cannot be perfectly negatively correlated to the other two; if ρ_{xy} and ρ_{xz} are specified as -1, ρ_{yz} is by consequence equal to $+1$. Weaker constraints on this 'consequential correlation' ρ_{yz} exist when the specified correlations are intermediate. Applying the method described herein to a 3-patch (or n-patch) system would require only a modification of the 'coin toss' procedure outlined in the methods; by specifying ρ_{xy} and ρ_{xz}, and by correlating the sequence of 'coin tosses' used to define their phase shifts ρ_{yz} can be varied across the constrained range of correlations. Furthermore, although the examples shown here are limited to the cases where the spectral exponent γ does not vary among local environments, this is not a requirement of the method or of nature (albeit heterogeneous spectral exponents impose further restrictions on the correlation between patches). Coastal and continental terrestrial habitats are known to possess quite different spectral exponents γ (Pelletier 2002, Vasseur and Yodzis 2004), and for aquatic environments, the spectral exponent can be closely related to the size of the water body (Cyr and Cyr 2003). Further research into the impact of varied patch arrangements and dispersal on population synchrony in coloured environments will supplement our understanding of how synchrony is achieved across large spatial scales and environmental gradients.

Acknowledgements – I would like to thank Andrew Gonzalez and Jeremy Fox for commenting on an early version of the manuscript. Jörgen Ripa provided helpful discussion of the method. Mike Fowler and Walter Koenig provided valuable feedback during the review process. Funding for this project was provided by post-doctoral fellowships from the Natural Sciences and Engineering Research Council of Canada and from Alberta Ingenuity.

References

Abbott, K. C. 2007. Does the pattern of population synchrony through space reveal if the Moran effect is acting? – Oikos 116: 903–912.

Bjørnstad, O. N. et al. 1999. Spatial population dynamics: analyzing patterns and processes of population synchrony. – Trends Ecol. Evol. 14: 427–432.

Chatfield, C. 2004. The analysis of time series: an introduction. – CRC Press.

Cohen, J. E. et al. 1999. Spectral mimicry: a method of synthesizing matching time series with different Fourier spectra. – Circuits Systems Signal Process. 18: 431–442.

Cuddington, K. M. and Yodzis, P. 1999. Black noise and population persistence. – Proc. R. Soc. Lond. B 266: 969–973.

Cyr, H. and Cyr, I. 2003. Temporal scaling of temperature variability from land to oceans. – Evol. Ecol. Res. 5: 1183–1197.

Engen, S. and Sæther, B.-E. 2005. Generalizations of the Moran effect explaining spatial synchrony in population fluctuations. – Am. Nat. 166: 603–612.

Fontaine, C. and Gonzalez, A. 2005. Population synchrony induced by resource fluctuations and dispersal in an aquatic microcosm. – Ecology 86: 1463–1471.

Greenman, J. V. and Benton, T. G. 2001. The impact of stochasticity on the behaviour of nonlinear population models: synchrony and the Moran effect. – Oikos 93: 343–351.

Grenfell, B. T. et al. 1998. Noise and determinism in synchronized sheep dynamics. – Nature 394: 674–677.

Halley, J. M. 1996. Ecology, evolution and 1/f-noise. – Trends Ecol. Evol. 11: 33–37.

Halley, J. M. 2004. Uses and abuses of fractal methodology in ecology. – Ecol. Lett. 7: 254–271.

Heino, M. 1998. Noise colour, synchrony, and extinctions in spatially structured populations. – Oikos 83: 368–375.

Hudson, P. J. and Cattadori, I. M. 1999. The Moran effect: a cause of population synchrony. – Trends Ecol. Evol. 14: 1–2.

Hugueny, B. 2006. Spatial synchrony in population fluctuations: extending the Moran theorem to cope with spatially heterogeneous dynamics. – Oikos 115: 3–14.

Ims, R. A. and Andreassen, H. P. 2000. Spatial synchronization of vole population dynamics by predatory birds. – Nature 408: 194–196.

Jeffrey, A. 2004. Handbook of mathematical formulas and integrals, 3rd ed. – Elsevier Academic Press.

Kaitala, V. et al. 1997. Population dynamics and the colour of environmental noise. – Proc. R. Soc. Lond. B 264: 943–948.

Kendall, B. E. et al. 2000. Dispersal, environmental correlation, and spatial synchrony in population dynamics. – Am. Nat. 155: 628–636.

Koenig, W. D. 2001. Spatial autocorrelation and local disappearances in wintering North American birds. – Ecology 82:2636–2644.

Koenig, W. D. 2002. Global patterns of environmental synchrony and the Moran effect. – Ecography 25: 283–288.

Koscielny-Bunde 1998

Koscielny-Bunde, E. et al. 1998. Indication of a universal persistence law governing atmospheric variability. – Phys. Rev. Lett. 81: 729–732.

Laakso, J. et al. 2003. Environmental noise and population dynamics of the ciliated protozoa *Tetrahymena thermophila* in aquatic microcosms. – Oikos 102: 663–671.

Lande, R. et al. 1999. Spatial scale of population synchrony: environmental correlation versus dispersal and density regulation. – Am. Nat. 154: 271–281.

Liebhold, A. et al. 2004. Spatial synchrony in population dynamics. – Annu. Rev. Ecol. Evol. Syst. 35: 467–490.

Liebhold, A. et al. 2006. Geographic variation in density-dependent dynamics impacts the synchronizing effect of dispersal and regional stochasticity. – Popul. Ecol. 48: 131–138.

May, R. M. 1976. Theoretical ecology, principles and applications. – Blackwell.

Moran, P. A. P. 1953. The statistical analysis of the Canadian lynx cycle. II. Synchronization and meteorology. – Aust. J. Zool. 1: 291–298.

Pelletier, J. D. 2002. Natural variability of atmospheric temperatures and geomagnetic intensity over a wide range of time scales. – Proc. Natl Acad. Sci. USA 99 (suppl. 1): 2546–2553.

Peltonen, M. et al. 2002. Spatial synchrony in forest insect outbreaks: roles of regional stochasticity and dispersal. – Ecology 83: 3120–3129.

Petchey, O. L. 2000. Environmental colour affects aspects of single-species population dynamics. – Proc. R. Soc. Lond. B 267: 747–754.

Petchey, O. L. et al. 1997. Effects on population persistence: the interaction between environmental noise colour, intraspecific competition and space. – Proc. R. Soc. Lond. B 264: 1841–1847.

Ranta, E. et al. 1995. Synchrony in population dynamics. – Proc. R. Soc. Lond. B 262: 113–118.

Ranta, E. et al. 1997. The Moran effect and synchrony in population dynamics. – Oikos 78: 136–142.

Ranta, E. et al. 1999. Spatially autocorrelated disturbances and patterns in population synchrony. – Proc. R. Soc. Lond. B 266: 1851–1856.

Ranta, E. et al. 2006. Ecology of populations. – Cambridge Univ. Press.

Ripa, J. 2000. Analysing the Moran effect and dispersal: their significance and interaction in synchronous population dynamics. – Oikos 89: 175–187.

Ripa, J. and Lundberg, P. 1996. Noise color and the risk of population extinctions. – Proc. R. Soc. Lond. B 263: 1751–1753.

Ripa, J. and Ives, A. R. 2003. Food web dynamics in correlated and autocorrelated environments. – Oikos 64: 369–384.

Ripa, J. and Ranta, E. 2007. Biological filtering of correlated environments: towards a generalized Moran theorem. – Oikos 116: 783–792.

Ripley, B. D. 2006. Stochastic simulation. – Wiley.

Roughgarden, J. 1975. A simple model for population dynamics in stochastic environments. – Am. Nat. 109: 713–736.

Royama, T. 1992. Analytical population dynamics. – Chapman & Hall.

Royama, T. 2005. Moran effect on nonlinear population processes. – Ecol. Monogr. 75: 277–293.

Schwager, M. et al. 2006. Does red noise increase or decrease extinction risk? Single extreme events versus series of unfavorable conditions. – Am. Nat. 167: 879–888.

Small, R. J. et al. 1993. Synchronous and nonsynchronous population fluctuations of some predators and their prey in central Sweden. – Ecography 16: 360–364.

Tedesco, P. A. et al. 2004. Spatial synchrony in population dynamics of West African fishes: a demonstration of an intraspecific and interspecific Moran effect. – J. Anim. Ecol. 73: 693–705.

Vasseur, D. A. and Yodzis, P. 2004. The color of environmental noise. – Ecology 85: 1146–1152.

Weisstein, E. W. 2006. Sine – from Wolfram MathWorld. [WWW document] URL http://mathworld.wolfram.com/sine.html

Wichmann, M.C. et al. 2005. Extinction risk, coloured noise and the scaling of variance. – Theor. Popul. Biol. 68: 29–40.

Appendix 1.

Derivation of the necessary phase-shift operator δ

Since the phase shift operator δ is a constant applied to the reference phase θ_1 at every frequency f in Eq. 1, the desired cross correlation ρ_ξ is applied equally to every frequency contributing to Eq. 1. Therefore, it is sufficient to derive the relationship between ρ_ξ and δ for any single frequency on the interval $0 < f \leq n/2$.

The time series of x_1 and x_2 represent the dynamics of ξ_1 and ξ_2 at a single frequency f and are given by:

$$x_{1,t} = \sin(2\pi ft/n)$$
$$x_{2,t} = \sin(2\pi ft/n + \delta) \tag{A1}$$

Provided that n is reasonably large relative to f, $\bar{x}_1 = \bar{x}_2 = 0$ and

$$\sigma_{x_1} = \left[\frac{1}{n}\sum_{t=1}^{n}\sin^2(2\pi ft/n)\right]^{1/2} \tag{A2}$$

Given that summation of $\sin^2$ can be done in closed form by

$$\sum_{t=0}^{n}\sin^2(tc) = \left\{1 + 2n - \frac{\sin[c(1+2n)]}{\sin c}\right\} \Big/ 4 \tag{A3}$$

(Weisstein 2006), Eq. A2 can be written

$$\sigma_{x_1} = \left\langle \frac{1}{4n}\left\{1 + 2n - \frac{\sin[2\pi f(1+2n)/n]}{\sin[(2\pi f/n)]}\right\}\right\rangle^{1/2} \tag{A4}$$

Provided that n is large σ_{x_1} can be approximated by taking the limit as $n \to \infty$:

$$\lim_{n\to\infty}\sigma_{x_1} = \left\langle \frac{1}{4}\left\{\frac{1}{n} + 2 - \frac{\sin[2\pi f(1+2n)/n]}{n\sin(2\pi f/n)}\right\}\right\rangle^{1/2}$$
$$= 1/\sqrt{2} \tag{A5}$$

It can be easily shown using the same logic that $\sigma_{x_2} \to 1/\sqrt{2}$ for large n. Given Eq. A1 and A5 it follows that the expression for the cross-correlation of x_1 and x_2, given sufficiently large n is:

$$\rho_x = \frac{\frac{1}{n}\sum_{t=1}^{n}[\sin(2\pi ft/n)\sin(2\pi ft/n + \delta)]}{1/\sqrt{2}\cdot 1/\sqrt{2}} \tag{A6}$$

Using the Werner trigonometric product identity (Jeffrey 2004) Eq. A6 reduces to:

$$\rho_x = \frac{2}{n}\sum_{t=1}^{n}\left[\frac{\cos(-\delta) - \cos(4\pi ft/n + \delta)}{2}\right]$$
$$= \cos(-\delta) - \frac{1}{n}\sum_{t=1}^{n}\cos(4\pi ft/n + \delta) \tag{A7a, b}$$

For sufficiently large n the second term of Eq. A7b $\to 0$ and it follows that the necessary phase shift to generate a correlation ρ_x between two sinusoidal time series is

$$\delta = \cos^{-1}(\rho_x) \tag{A8}$$