Oikos

OIK-00819

Hamel, S., Yoccoz, N. G. and Gaillard, J.-M. 2013. A standardized approach to estimate life history tradeoffs in evolutionary ecology. – Oikos 000: 000–000.

Appendix A1

There is a mathematical constraint on the maximum negative correlation that can exist in binary data. The maximum correlation is constrained by the mean probability (Prentice 1988, Chaganty and Harry 2006) and by the standard deviation representing heterogeneity among individuals (Hamel et al. 2012). This constraint is illustrated in Fig. A1, where the maximum ρ (ρMAX) that can be observed has been determined through simulations (see Hamel et al. 2012 for how to determine ρMAX). The constraint does not occur for positive correlations (Hamel et al. 2012).

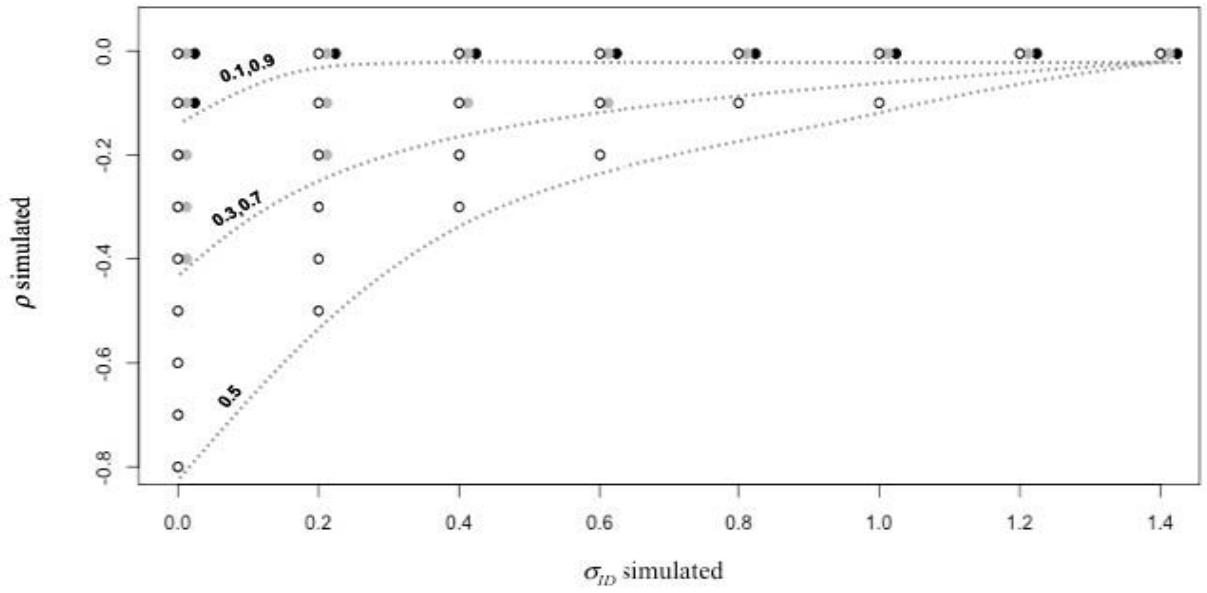

Figure A1. Combinations of autocorrelation and standard deviation that are possible for data following a Bernoulli process, according to different probabilities (open: 0.5, grey: 0.3 and 0.7, black: 0.1 and 0.9). Points were jittered. Dashed lines represent the lower theoretical limit for standard deviation and autocorrelation to occur at the same time according to each probability, and hence we should expect to find real population values above these limits. Reproduced from Hamel et al. (2012).

References

Chaganty, N. R. and J. Harry. 2006. Range of correlation matrices for dependent Bernoulli random variables. – Biometrika 93:197–206.

Hamel, S. et al. 2012. Statistical evaluation of parameters estimating autocorrelation and individual heterogeneity in longitudinal studies. – Meth. Ecol. Evol. 3:731–742.

Prentice, R. L. 1988. Correlated binary regression with covariates specific to each binary observation. – Biometrics 44:1033–1048.

Appendix A2

When estimating a tradeoff that includes lifetime traits, the correlation estimated is the relationship between the two traits for the entire population. When the tradeoff includes a trait that is recurrent during the lifetime, however, we have access to repeated measurements for the trait, for example at each age. We can therefore evaluate the autocorrelation between the sequential residuals with age from an expected trajectory. Because we have access to repeated measures, we can include a random effect that accounts for the variation among individual trajectories. If this random effect is an intercept, individuals are fitted with different intercepts but with the same slope (Fig. A2), whereas including a random slope also allows fitting different slopes for individuals. Therefore, in the case of recurrent tradeoffs, the autocorrelation is estimated among the sequential residuals around the expected trajectory of each individual.

In Fig. A2, we have illustrated a case where a random intercept is fitted. In A we show the variation among individuals, where each line represents the expected trajectory that is fitted to each individual (these fitted individual trajectories are represented by the dotted lines in B, C, D and E in a situation where the shrinkage of best linear unbiased predictors (BLUP) is negligible). If the life-history trait is body mass, then A could represent a situation where heavy individuals are heavier than lighter individuals at all ages. When using a time series model and assessing the autocorrelation with age, if the sequential residuals for an individual are independent then no correlation is found (case A). If a strong negative correlation is found, it indicates that the sequential residuals for an individual are oscillating between sides of the individual trajectory at each age (case B). If a positive correlation is found, then the residuals present sequences of positive or negative values (case C). Nevertheless, because the residuals are estimated from the expected trajectory of each individual, the residuals cannot always be above or below the individual trajectory. They instead necessarily oscillate around the trajectory, but over a long time (cases C, D and E). This oscillation over longer time frame can result from a biological effect or be a statistical artifact.

In the biological case, the positive correlation represents a sort of 'long-term memory'. For example, the influence of environmental conditions on individual size could take time to be observed in skeleton measurements, and we could find a positive correlation for skeleton measurements representing a long-term memory of the previous states (case C). On the other hand, the positive correlation could also be the result of an

inappropriate modeling of the age effect (cases D and E; Wakefield 2007, Hodges and Reich 2010). For example, if different trajectories going from fast-growing animals reaching high mass to slow-growing animals achieving only a small mass were present in the population (case D), one would need to fit a random slope model to account for the differences in growth pattern among individuals. Furthermore, if age is modeled as a linear effect but has a biological quadratic relationship, then the sequential residuals will be spuriously positively correlated (case E). The difference between case C and case E is that in E all individuals have the same quadratic curve with age that is inappropriately modeled linearly, resulting in a spurious positive correlation among the residuals that would disappear if a quadratic age effect was modeled. In C, however, adding a quadratic age effect would still leave a positive correlation among the residuals.

In roe deer, for instance, variation in body mass usually follows a quadratic relationship with age (Nussey et al. 2011). When we modeled the variation in sequential mass measurements for roe deer using the time series model, we first used a quadratic effect of age and found a statistically significant positive correlation of 0.16. To ensure this positive correlation corresponded to a real biological effect, we tried fitting different age effects. We first fitted a random slope model for age, with no influence on the results. However, when we then modeled age as a non-linear relationship in a generalized additive mixed model, the positive correlation disappeared and we found a correlation of 0 (Results). This illustrates the risk of finding spurious positive correlation as a result of inappropriate modeling of the fixed effects (Wakefield 2007, Hodges and Reich 2010), and we therefore strongly recommend paying extra attention to positive correlations to ensure that they are really the result of a biological difference and do not correspond to a statistical artifact.

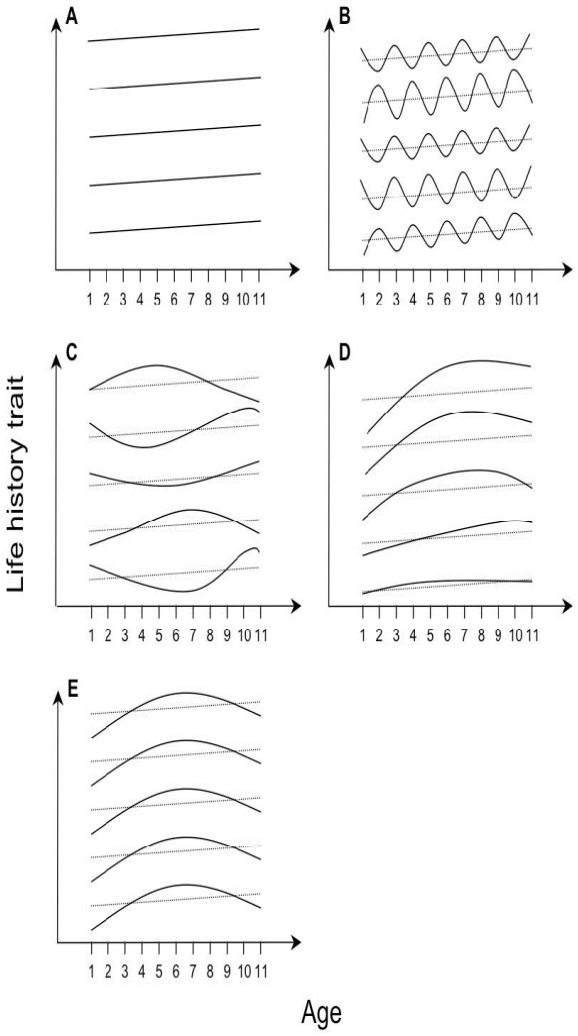

Figure A2. Illustration of different combinations of individual variation in a life-history trait with age. The continuous lines represent the actual individual trajectories; the dotted lines represent the best linear unbiased predictors (BLUP), in a situation where the shrinkage is negligible.

References

Hodges, J. S. and Reich, B. J. 2010. Adding spatially-correlated errors can mess up the fixed effect you love. – Am. Stat. 64: 325–334.

Nussey, D. H. et al. 2011. Patterns of body mass senescence and selective disappearance differ among three species of free-living ungulate. – Ecology 92:1936–1947.

Wakefield, J. 2007. Disease mapping and spatial regression with count data. – Biostatistics 8:158–183.

Appendix A3

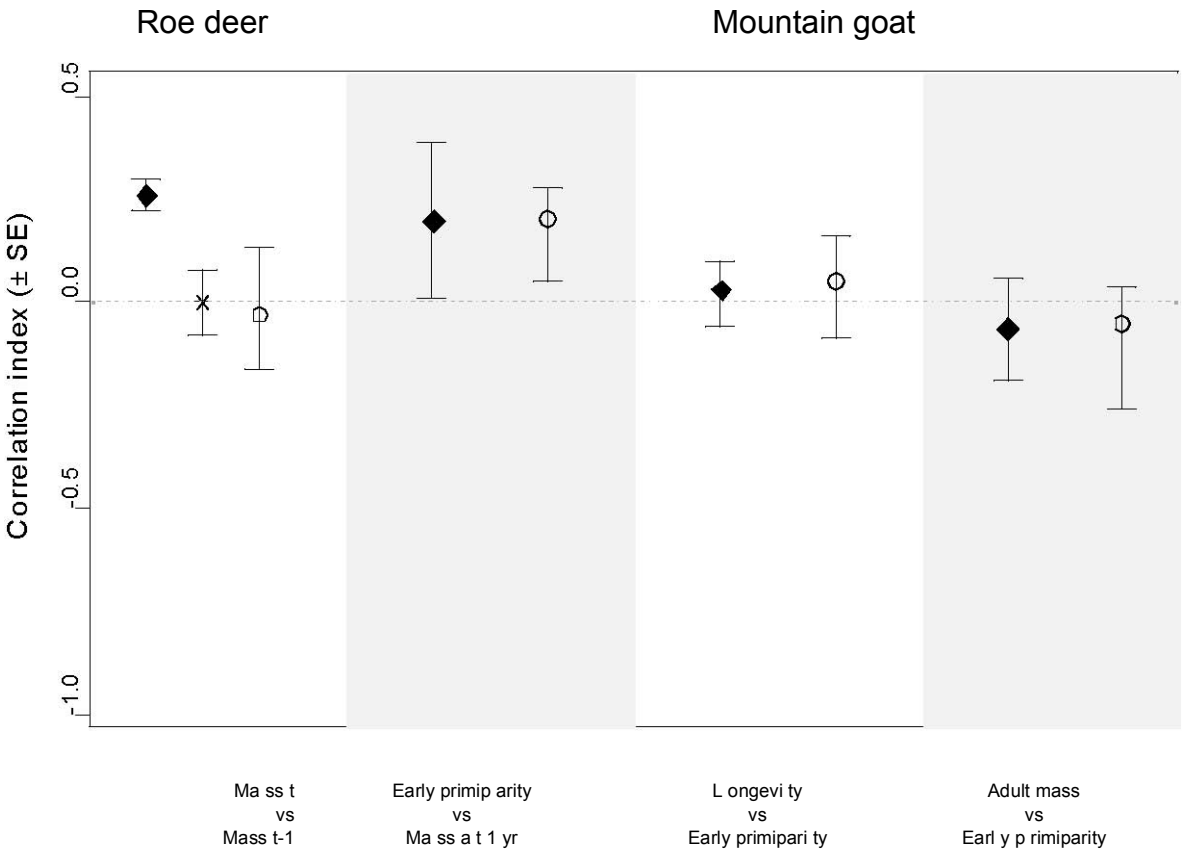

Figure A3. Comparison of the correlation coefficients obtained using the covariance estimated from MCMCglmm models (open circles; Hadfield 2010) versus estimates from the state dependence models (i.e. βT-values; diamonds) and the time series models (i.e. ρ-values; stars). The correlation coefficient represents the strength of tradeoffs among different life history traits for roe deer and mountain goats. Whiskers represent standard errors. The time series model only provides tradeoff estimates for the first traits because this model can only estimate ρ for a tradeoff occurring within a single trait measured at successive periods of the life cycle. We cannot compare the MCMCglmm method for the binary traits because the covariance estimated by this model for binary data will always be 0 (see text).

References

Hadfield, J. D. 2010. MCMC methods for multi-response generalized linear mixed models: the MCMCglmm R package. – J. Stat. Software 33: 1–22.

Appendix A4

Some life history traits follow distributions not well approximated by the normal or Bernoulli/binary distributions. Indeed, another common distribution of life history traits is the Poisson distribution (Fujiwara 2007). A Poisson distribution arises from a trait being measured as counts of occurrences and is characterized by a single parameter, λ, which is equal to both the mean and the variance of the trait. Here, we demonstrate using simulations that the first-order autocorrelations estimated from the time series models, ρ, are similar to the βT coefficients estimated from the state dependence models for traits following a Poisson distribution, and that both estimates can be used to quantify life history tradeoffs.

For the normally distributed traits, the response and explanatory variables must be standardized before running the models to allow comparisons among βT, as well as between βT and ρ. Because of the nature of the Poisson distribution, however, the models must be run on data that are positive integers, thereby making it irrelevant to standardize the data before the analyses. Therefore, the standardization for βT needs to be done a posteriori, by multiplying the βT estimated with the variance of the trait, i.e. with λ (see the R script at the end of the appendix).

We simulated autocorrelated Poisson data using the NORTA method (NORmal To Anything; Chen 2001) that we have adapted to simulate time series data (see R script below). The principle behind the NORTA method is simply to generate first normally distributed data with the desired correlation structure and then to transform it into the distribution we wish to simulate (in our case Poisson) using the inverse cumulative distribution function (Chen 2001, Yakav and Shmueli 2012). For a Poisson distribution with low λ, however, the actual correlation simulated is often lower than the desired correlation, and so a correction must be applied to simulate the appropriate correlation (Yakav and Shmueli 2012). We therefore applied Yakav and Shmueli (2012) correction to simulate the Poisson data (see the R script below). We simulated datasets of 100 individuals living for 20 years for different combinations of negative autocorrelation (ρ: 0, –0.2, –0.4, –0.6, –0.8) and of mean/variance of traits (λ: 2, 4, 10). We then used the R function glmmPQL (Venables and Ripley 2002) to run both the time series and the state-dependence models and evaluated the ρ and the βT estimated, respectively. We then standardized βT a posteriori by multiply it with λ. We also evaluated the βT estimated by the state-dependence model using the R function glmer (Bates et al. 2012). Since the βT obtained from glmmPQL and glmer were the same, we only presented results for the former.

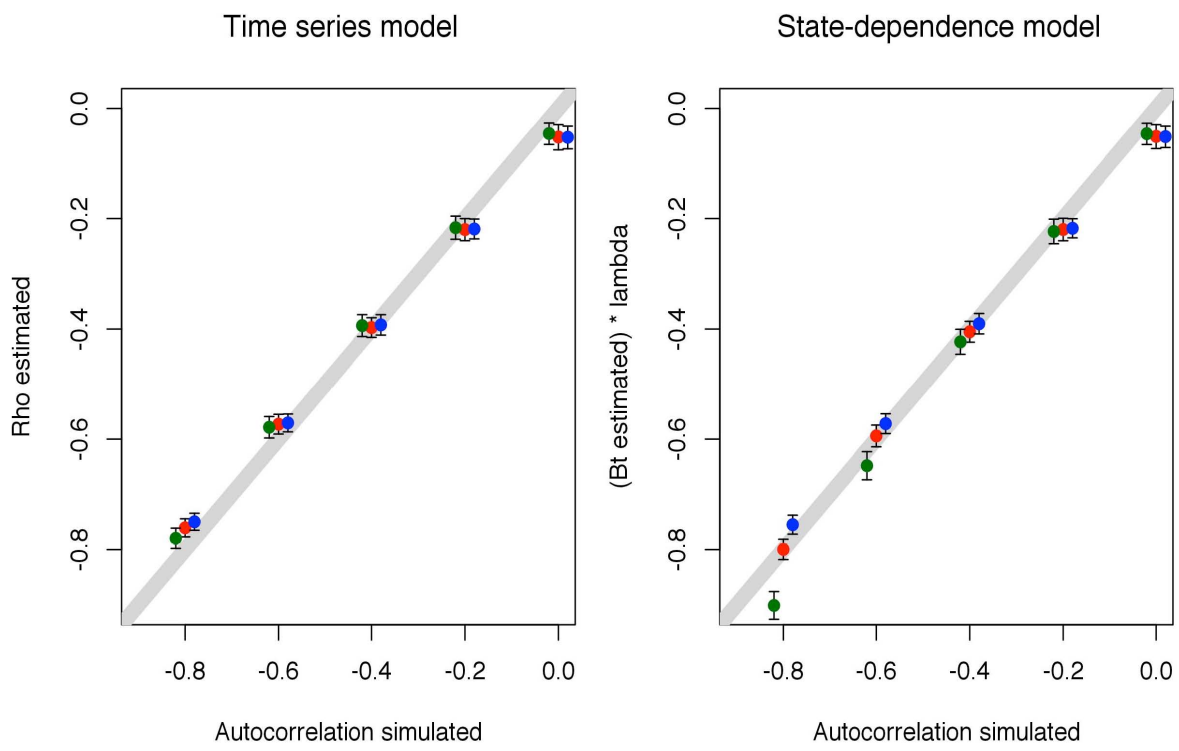

Figure A4. Comparison of the negative autocorrelation values simulated with the autocorrelation (i.e. ρ) estimated from the time series model (left panel) and with the correlation (i.e. $\beta_T \times \lambda$) estimated from the state-dependence model (right panel), based on Poisson distributed data. Estimates represented means ± standard deviations from 100 simulations, and are presented for three different values of λ (green: 2, red: 4, blue: 10). Estimates have been jittered on the x-axis to facilitate visualisation, and the grey diagonal represents the perfect fitting line between values simulated and estimated.

R script used to simulate autocorrelated time series with Poisson distribution:

```
require(compiler);require(MASS);require(lme4)
enableJIT(3)

# Function to correct initial correlation between a certain pair of series (following Yahav and Shmueli 2012).


# lambda1: lambda of first year (or age)
# lambda2: lambda of second year (or age)
# r: desired correlation to be simulated

correct.init.corr = function(lambda1, lambda2, r){
        samples=500
        u = runif(samples,0,1)
        lambda=c(lambda1,lambda2)
        maxcor=cor(qpois(u, lambda1), qpois(u, lambda2))
        mincor=cor(qpois(u, lambda1), qpois(1-u, lambda2))
        a=-maxcor*mincor/(maxcor+mincor)
        b=log((maxcor+a)/a, exp(1))
        c=-a
        corrected=log((r+a)/a, exp(1))/b
        corrected= ifelse ((corrected>1 | corrected<(-1)), NA, corrected)
        return(corrected)

}

# Function to simulate correlated random Poisson time series data, using the NORTA method adapted for time
series simulation

# Nage: length of the age time series to simulate
# Nind: number of individual to simulate
# lambda: the mean and variance to simulate for the trait studied -we use the same as it is the same repeated trait in
this      case #
r: desired correlation to be simulated

poisson.serie.corr = function(Nage, Nind, lambda, r) {
        rho = correct.init.corr(lambda,lambda,r)
        yt1=yt=matrix(NA,nrow=Nage-1, ncol=Nind)
        for (i in 1:Nind){
                normal = arima.sim(model=list(ar=rho), n=Nage)
                normalscaled = scale(normal)
                unif = pnorm(normalscaled)
                y = t(qpois(t(unif), lambda))
                yt1[1:(Nage-1),i]=tail(y,-1)
                yt[1:(Nage-1),i]=head(y,-1)
        }

        yt1 = as.vector(yt1)
        yt = as.vector(yt)
        x = (rep(c(1:(Nage-1)),times=Nind))
        ran = rep(c(1:Nind), each=(Nage-1))
        obs = 1:length(yt1)
        data=data.frame(yt1=yt1,yt=yt,x=x,ran=ran,obs=obs)
        return(data)
}

# Function to run the time series model and extract the rho estimate
cor.ar1=function(dataset){
        glmm1=try(glmmPQL(fixed=yt1~x, random=~1|ran, family=poisson, correlation=corAR1(-0.4,
```

```
form=~1|ran),     data=dataset))
        if (inherits(glmm1, "try-error")) cor=NA
        if (!inherits(glmm1, "try-error")) cor=as.vector(coef(glmm1$modelStruct$corStruct,
unconstrained=FALSE))    return(cor)
        }

# Function to run the state-dependence model, extract the βT estimate and standardize it with λ (using glmmPQL)
cor.btpql=function(dataset){
        glmm1=try(glmmPQL(fixed=yt1~yt+x, random=~1|ran, family=poisson, data=dataset))
        correction=var(dataset$yt1)
        if (inherits(glmm1, "try-error")) cor=NA
        if (!inherits(glmm1, "try-error")) cor=((as.vector(fixef(glmm1)))[2])*correction
        return(cor)
        }
# Function to run the state-dependence model, extract the βT estimate and standardize it with λ (using glmer and
        controlling for potential overdispersion by adding an individual-level random effect – i.e. variable "obs",
see     Agresti 2002 section 13.5)
cor.btpql=function(
cor.btlmer=function(dataset){
        glmm1=try(glmer(yt1~yt+x + (1|ran)+ (1|obs), family=poisson, data=dataset))
        correction=var(dataset$yt1)
        if (inherits(glmm1, "try-error")) cor=NA
        if (!inherits(glmm1, "try-error")) cor=(as.vector(glmm1@fixef[2]))*correction
        return(cor)
        }

rholist=list(0,-0.2,-0.4,-0.6,-0.8) # List of autocorrelation values desired to be simulated
nsim=100 # Number of simulations wanted

list.datasets=replicate(nsim,lapply(rholist, poisson.serie.corr, Nage=20, Nind=100, lambda=4))
lapply(list.datasets, cor.ar1)
lapply(list.datasets, cor.btpql)
lapply(list.datasets, cor.btlmer)
```

References

Agresti, A. 2002. Categorical data analysis. – Wiley.

Bates, D. et al. 2012. lme4: Linear mixed-effects models using S4 classes.

Chen, H., 2001. Initialization for NORTA: generation of random vectors with specified marginals and correlations. – INFORMS J. Comp. 13: 312–331.

Devenish-Nelson, E. S. et al. 2013. Does litter size variation affect models of terrestrial carnivore extinction risk and management? – PLoS ONE 8, e58060.

Fujiwara, M. 2007. Extinction-effective population index: incorporating life-history variations in population viability analysis. – Ecology 88: 2345–2353.

Kendall, B. E. and Wittmann, M. E. 2010. A stochastic model for annual reproductive success. – Am. Nat. 175: 461–468.

Venables, W. N. and Ripley, B. D. 2002. Modern applied statistics with S. – Springer.

Yakav, I. and Shmueli, G. 2012. On generating multivariate Poisson data in management science applications. – Appl. Stochastic Models Bus. Ind. 28: 91–102.