

Evaluation of the statistical analyses in Hallmann et al. 2014 “Declines in insectivorous birds are associated with high neonicotinoid concentrations”.

5.1.2.e Wageningen University and Research Centre
14 August 2014, Wageningen. 5.1.2.e [@wur.nl](mailto:5.1.2.e@wur.nl)

Evaluatie van de statistische analyses in Hallmann et al 2014 “Declines in insectivorous birds are associated with high neonicotinoid concentrations.

Korte samenvatting

De statistische analyses in Hallmann et al (Nature, 2014) worden kritisch bekeken vanuit statistisch perspectief. De zorgen vanuit dit perspectief blijken na analyse van de samengevatte gegevens niet tot andere conclusies te leiden.

Eén conclusie moet wel naar mijn mening wel worden afgezwakt. Het is de conclusie die getrokken wordt uit de vergelijking van trendschattingen uit twee perioden (1984-1995 en 2003-2010) en hun verband met gemiddelde log imidacloprid concentratie uit de periode 2003-2009. Het verschil in effect tussen die twee periodes is vermoedelijk niet statistisch significant.

Het artikel heeft een goede discussie over de vraag of de correlatie van afname met imidacloprid op een causaal verband kan wijzen.

Een aantal aspecten kan alleen worden beantwoord in een uitgebreide heranalyse van de originele gegevens. Het rapport bevat een aantal aanknopingspunten hiervoor.

Uitgebreide samenvatting

In Hallmann et al (Nature, 2014) wordt de afname in insectenetende vogels in verband gebracht met hoge concentraties imidacloprid (imid) op basis van data uit gestandaardiseerde monitoring data van Sovon en de bestrijdingsmiddelenatlas van Nederland. Dit wordt gedaan op basis van statistische analyses. In dit rapport worden de statistische analyses in het artikel kritisch bekeken vanuit statistisch perspectief. De samengevatte gegevens die met het artikel gepubliceerd zijn (weergegeven in Figuur 1b van het artikel), zijn op verschillende manieren nader geanalyseerd.

Het belangrijkste punt van zorg over de gebruikte analyses is dat ze lijden aan pseudo-herhaling. Bij pseudo herhaling worden gegevens als onafhankelijk beschouwd terwijl ze feitelijk statistisch gezien afhankelijk zijn. Zo worden de trendschattingen van soorten van eenzelfde monitoring plot als onafhankelijke punten in de analyse opgenomen en wordt er geen rekening gehouden met mogelijke ruimtelijke afhankelijkheid (autocorrelatie) en er wordt ook niet achteraf gecontroleerd of autocorrelatie een rol kan spelen.

Een punt van zorg is de interpolatie van imid naar de monitoring plots. Daarbij is gebruik gemaakt van kriging met externe drift. De gebruikte drift vorm kan ongewenste randeffecten geven. De mogelijke randeffecten (voor zover leidend tot hoge concentraties) lijken echter het verband niet direct te treffen aangezien de trendschattingen in de hoogste concentratie klasse juist hoger zijn dan verwacht op basis van het lineaire verband met log-imid. De onnauwkeurigheid in de imid

metingen en interpolaties zijn geen punt van grote zorg aangezien het niet leidt tot het overschatten van het verband; het leidt eerder tot onderschatting van de sterkte.

Trendschattingen uit twee perioden (1984-1995 en 2003-2010) worden in verband gebracht met de gemiddelde log-imid concentratie uit de periode 2003-2009. Voor de periode 1984-1995 wordt geen statistisch significant verband gevonden en voor de periode 2003-2010 wel. Het verschil tussen “statistisch significant” en “statistisch significant” is op zich niet statistisch significant. Het verschil is mogelijk niet significant hetgeen de onderbouwing ondermijnt van de conclusies die uit dit verschil worden getrokken in het artikel.

Een sterk punt van het artikel is het onderzoek naar acht andere mogelijke verklaringen voor de afname in insectenetende vogels. In die vergelijking is imid veruit de belangrijkste factor en zijn zeven van de acht variabelen niet statistisch significant. Het probleem van pseudo-herhaling in de gegevens, zoals hierboven omschreven, is kennelijk niet zo groot dat erg gemakkelijk statistische significante verbanden gevonden kunnen worden.

De samengevatte gegevens leiden veel minder aan mogelijke pseudo-herhaling en hun aantal is kleiner dan het aantal onderzochte monitoring plots. De statistische analyses uitgevoerd op de samengevatte gegevens laten zien dat het statistische verband tussen de afname in insectenetende vogels en imid robuust is tegen grote perturbaties van de gegevens. Zo leidt het weglaten van de laagste twee imid klassen of de hoogste twee klassen (die mogelijk een effect van de imid interpolatie zijn) niet tot andere conclusies. Zelfs op basis van maar negen datapunten, die gevormd zijn door de mediaan te nemen van de trendschattingen per imid klasse, is er statistisch gezien nog een duidelijke aanwijzing voor het verband.

In het artikel wordt zorgvuldig omgegaan met de vraag of een correlatie ook een aanwijzing voor een causale verband zou kunnen betekenen.

Een aantal aspecten kan alleen worden beantwoord in een uitgebreide her-analyse van de originele gegevens. Het rapport bevat een aantal aanknopingspunten hiervoor.

Introduction

This report discusses the statistical aspects of Hallman et al. (2014) upon a request by the Ctgb-College voor de toelating van gewasbeschermingsmiddelen en biociden on 30 July, 2014. The deadline for the report was 14 August, 2014. The original data were not available but summarized data turned out to be available.

References to the paper by Hallman et al. (2014) take the form N:... or N:Sp... for a reference to the main paper or the Supplementary Information page

The R code with results for the summarized data is added as Appendix to this report.

1. Issues when inferring from spatial patterns

The main issue when inferring relationships from two spatial pattern is that of spatial auto-correlation in either or both spatial variables. Spatial auto-correlation entails that the values at two close locations is likely to be more similar than those at two distant locations. Autocorrelation can be seen as a problem of pseudo-replication in that two very highly correlated observations should in fact be counted as one observation. This is the basis of, for example, Dutilleul's modified t-test (Dormann et al. 2007). The most striking feature is that the uncertainty of a regression coefficient increases. More detail is as follows.

In a regression setting as in the paper, most attention has be devoted to autocorrelation in the response variable (here the local bird trend) as standard regression results are based on the assumption of independence of the errors. This assumption can be partly checked by an analysis of the residuals of the fitted regression model. In the paper, it would have been logical to check for residual spatial autocorrelation by fitting a semivariogram to the residuals of the analysis in the same way as done in the paper for the imidacloprid (for short, imid) concentrations (Extended Data Figure 3). Such or similar analysis has not been reported. If residual autocorrelation is found (or is suspected), a method to account for spatial autocorrelation can be applied (Dormann et al. 2007). The effect of auto-correlation can give bias in the estimate of the effect in some situations, but not necessarily so as Dormann et al (2007) confirmed in simulations. "The more striking feature is underestimation of the uncertainty." (Dormann et al. 2007).

Less attention has been devoted to autocorrelation in the predictor variable (here log imid concentration). In time series analysis, which is a slightly easier context, the standard used to be to pre-whiten the two time series before correlating them. Pre-whitening is an operation that attempts to remove the autocorrelation in a time series. The standard model nowadays is the Vector AutoRegression (VAR) model (http://en.wikipedia.org/wiki/Vector_autoregression) which does the whitening simultaneously with estimation of the relation between the two series.

High (spatial) autocorrelation in the predictor makes it more likely that there is another (unobserved) variable that may also explain the response variable. If, for example, imid takes only two values (low and high), and low occurs only in the East of the Netherlands and high only in the West of the Netherlands, than any variable that differs greatly between East and West can equally explain the bird trend. The problem of spatial autocorrelation is thus closely linked to the possibility of other variables causing the effect on bird trends.

In N:Figure 1a there appear to be 8-12 disjoint regions with concentrations >50ng/l and 2-4 disjoint regions with concentrations <10ng/l. For a regression analysis, as reported in N:Figure 1b, the low regions are equally informative as the high regions.

2. Interpolated and imprecise imid values

The imid concentrations are low, are measured with error and all water boards have their own measuring method. These issues are largely outside the scope of this evaluation but I have a few general remarks on the possible influence of measurement error and of the interpolation.

Measurement error

The statistics literature distinguishes two main types of measurement error (ME), classical and Berkson. For a recent overview see Muff et al. (2014). In the classical error model, the predictor entered in the regression has a higher variance than the true covariate whereas in the Berkson case the opposite applies. Their effects on the estimated regression coefficient differ: with classical ME, the regression coefficient is attenuated, meaning that it has a bias toward 0 (http://en.wikipedia.org/wiki/Errors-in-variables_models), whereas with Berkson ME the regression coefficient is unbiased. Both types of ME thus imply that the effect in the paper is not overestimated due to ME. If imid would have been measured at the monitoring plot, ME would certainly have been of the classical type and the effect of imid is likely to be underestimated.

Interpolation by kriging

The imid concentration were not available at the monitoring plots and were interpolated by universal kriging with external drift (a third order polynomial in the geographic coordinates). The sensitivity of the results to changes in this procedure has not been reported.

The third order polynomial drift is a drift that is likely to cause edge effects. The high interpolated values in East Groningen may perhaps be explained by the high values at neighboring plots in North and South-East Groningen, but the high interpolated values in East Twente (>50ng/l) are likely due to this drift specification. The influence might be that some of the points in the highest classes of imid in N:Figure 1b move to the left. This is not likely to change much the estimated effect.

Interpolation by kriging can also be a good thing in that the procedure may smooth the variability in imid and thus also lessen the variability due to measurement error.

In theory, the problem that the imid concentration were not available at the monitoring plots, is one of missing values. A general statistical methodology for missing data is that of multiple imputation (Rubin 1996, van Buuren 2012). By using this methodology the uncertainty caused by the interpolation could have been accounted for. The likely effect of the uncertainty is that the standard error of the regression coefficient would increase.

I recommend using the multiple imputation methodology in any full re-analysis of the data.

3. Detailed statistical evaluation

3.1 The data used

Bird data

The selection of bird species is on subject matter grounds (N:Sp1-4): insectivorous farmland passerine species that are common to the Dutch agricultural landscape. The national trend estimates over the period 1990-2005 are reported in qualitative terms in N:Extended Data Table 1: four species with increasing trends, four stable trends and seven negative trend in the period 1990-2005. This period is largely before the main (eight years) period studied in the paper (2003-2010).

Monitoring plots “with less than three years of non-zero counts for a given species were” ignored. This selection is irrespectively of the years in which the non-zero

count occurred and thus is not likely to be biased. A local trend is thus calculated for each plot on a time series of 10 years which has at least three non-zero counts. This is a bare minimum, which will lead to sometimes very imprecise trend estimates per plot. Imprecision will make the detection of a statistical relation of trends with imid more difficult. The method of trend estimation is a Poisson log-linear regression of territory counts against year (N:Sp8). The trend is estimated by the slope of this regression and this is the standard method.

Monitoring plots that are >5 km from imid measurement points are also ignored. All 354 remaining monitoring plots (17-200 for individual species) are thus within 5 km of imid measurement points “to avoid basing inference on distant interpolations”(N:Sp4). The locations of the 354 plots are shown in N:Extended Figures 2.

Imidacloprid data

The measurements are from the period 2003-2009 and carried out in March – September. Data with limits of concentration above 10 ng/l were ignored. In total 3947 measurements were used from 555 unique sampling locations. Per location, the log-concentrations are averaged and the resulting averages (z) are submitted to a geo-statistical analysis (universal kriging with drift) to map to a fine grid (250 meter) across the Netherlands represented as color map with contour lines in N:Fig.1a. The drift is a third order polynomial of the spatial coordinates which may lead to unwanted edge effects at points further away from sampling locations. It would have been an advance if the sensitivity to this choice was investigated and reported. The kriging used a “classical matern variogram” (N:Sp5) where “classical” refers here, as I understand it, to variogram and not matern. The roughness parameter (κ) of the Matérn model is not reported, but can influence the interpolation. The reported range of spatial autocorrelation is reported as 13.2 km with a nugget of 0.1901 and a sill of 1.6989. Nugget and sill are the variances of the averaged log-concentration (z) at very short and very large distances, respectively. The value of 0.19 translates to a coefficient of variation of about 44% of the geometric mean concentration ($\exp(z)$) (Jongman et al. 1995)[p.19]. At 5 km, the semi-variance is about 1 (N: Extended Data Figure 3), which translates to a coefficient of variation of more than 100% of the geometric mean concentration ($\exp(z)$). There is thus considerable variability in measured concentration, even at short distances. The effect of measurement error in covariates (predictors) of a regression is discussed in section “”Measurement error””.

The software used for kriging is the R library `gstat`. It is not reported whether the variogram estimation took into account the drift model. As it is not reported, it is likely that the variogram was calculated from the raw imid averages and not in a more advanced way, such as from residuals from the third order polynomial drift model or in a model-based way.

The temporal correlation of imid is found to be high in this data (N: Extended Data Figure 3). Vijver and van den Brink (2014) analyzed temporal variability of imid and found it to be high (implying low temporal correlation), but their analysis of variability does not use formal statistical tools.

The imid concentrations used in the paper are averages of the kriged log-concentrations at the grid-cells whose centroid fell within the monitoring plot boundaries. The resulting values are denoted by x in the R code in the Appendix to this report.

Kriging can also output the uncertainty estimate of a prediction (an interpolation), but this information is neither used nor reported nor mapped.

It would have been useful to investigate the sensitivity of the main statistical analysis on the choice of the drift model.

Further remarks on the data

I did not check whether the locations of the monitoring plots in N:Extended Data Figure 2 are all within 5 km of a imid sampling station in N:Extended Data Figure 1, but I wondered whether the locations in East Groningen and Twente fulfilled the criterion. From a statistical point the criterion has importance in view of the range of the autocorrelation. N:Extended Data Figure 5b shows that the 5 km is not crucial to the conclusions of the paper; in the range 1-25 km, the smaller the cut-off value the stronger is the negative effect of imid on bird trends.

A striking feature of N:Figure 1a is the above 100 ng/l concentrations in East-Groningen and in Twente. This feature may be an edge effect of the drift model, but note that in East Groningen sampling locations with high imid concentrations are nearby (N: Extended Data Figure 1).

There are no monitoring plots in N. Limburg or near Venlo, which is apparently a short-coming in the Dutch Common Breeding Monitoring Scheme. Such plots would have been useful as there is a high average imid concentration in this region.

3.2 Statistical analyses of bird trend on imid

The joint statistical analysis on all 15 species (LMER)

The main analysis reported in the main text of Hallmann et al. (2014) and summarized in N:Figure 1b is an analysis of all 15 species jointly (pooled analysis). The data are available from the Nature web site in the discretized form presented in N:Figure 1b (http://www.nature.com/nature/journal/v511/n7509/source_data/nature13531-f1.csv) The file also contains the species name attached to each point in the figure and thus allows to check the analysis albeit on binned data. Binning is not necessarily a bad thing as the interpolated imid data contain error anyway.

N:Figure 1b is based on a model with 15 parallel lines, one for each species, thus with a common slope and with an (random) intercept that depends on the particular species, of which the average line is shown with confidence bands. The model is fitted to the (averaged) interpolated imid log-concentrations. The model suffers in my opinion from pseudo-replication (Hurlbert 1984) in that it treats the trends of different species at the same sampling location as independent units (after correction for species-specific random intercept) whereas they are in fact dependent. Pseudo-replication underestimates variability and usually leads to statistical tests that are too tolerant, in particular, tests that judge an effect too often as significant where there is no true effect. The model treated the $n = 1459$ residuals of the model as independent with no check on the validity of this. My re-analyses try to avoid pseudo-replication as much as possible and are based on $n = 119$ or even on $n = 9$ observations. These analyses, reported in section “Re-analysis of the binned data”, show that effect of imid on bird trends, as deducible from the binned data, remains strong and significant.

The model with species-dependence slopes was run on the full data by the authors as reported in the N:Sp9 but not used “as the random slope to imidaclo-prid concentrations for each species accounted for an insignificant amount of variance (sd=0.0070) compared to the intercept (sd=0.0375),...”. The sd for slopes depends on the scale of log-imid concentrations and is thus difficult to compare with the sd of the intercept (which has a fixed scale). Although the written motivation for use of the simple model is not convincing it is convincing knowing the range of log-imid concentrations. Moreover, the choice for the simple model appears valid as judged from my re-analysis on the basis of the binned data. The simple model fitted as well as both the correlated and the uncorrelated random-intercept-random-slope model as judged by likelihood ratio tests using lme4 (Bates et al. 2014) (p-values>0.7). The differences in slopes among species are thus not statistically significant. Anyway, the conclusions about the fixed effect of imid on bird trends are about the same under the more complicated models. Note that these more complicated models suffer from the same pseudo-replication as the simple model.

I missed checks for residual correlation among species trends after fitting the model so as to be able to judge the amount of pseudo-replication. The simplest way of extending the model to allow for pseudo-replication is to add a random effect for monitoring plot as applied in a slightly different context by Jamil et al. (2013); even that extension may not be sufficient to remove all residual correlation.

I also missed a plot of the spatial auto-correlogram or of the semi-variogram of the residuals of the analysis, to check for residual spatial autocorrelation and (in connection with the previous paragraph) cross-correlations among species.

Interpretation of the joint analysis

The regression line $y = 0.111 - 0.0374 \log(\text{imid})$ in N:Figure 1b is used in the main text (N:p1) to derive that “the imid concentration above which bird populations were in decline was 19.43 (s.e.m 0.03) ng/l” with $19.43 = \exp(0.111/0.0374)$. In a model with random intercepts for species, some species are predicted to increase at this concentration, while other are predicted to decrease. Indeed, the average predicted trend across species is 0 at this concentration.

I note that the model is unrealistic in the sense that at low concentrations the species on average increase (only two decrease in my re-analysis) with more increase when extrapolating to even lower imid concentrations, whereas the overall trend in the Netherlands (in an earlier period) had a negative trend for 7 out of the 15 species. This aspect should have been discussed in the paper.

If the model holds in the range of the data, then the value of 19.43 ng/l has no special meaning in my opinion. If this model holds, the difference in trend at a particular concentration and the predicted trend at the lowest concentration or at the no-effect concentration is of interest. This reasoning makes the putative effects bigger than reported in the paper.

Comparison of the effects in two periods (1984-1995 and 2003-2009)

The trends per plot and species were also calculated for the period 1984-1995, i.e. before imid application and regressed to the imid data from the 2003-2009 period “in order to preclude that the interpolated imid concentrations reflect a spatial pattern of local farmland declines that developed before the introduction of imid in the Netherlands”. The calculated trends in this earlier period “did not relate to imid”, reported as for the full data as “(LMER: df 1288, $t = -1.56$, $P = 0.11$)” in N:Sp11 and for the monitoring plots with trends for both periods (thus using fewer plots) as ($t = -1.43$, $P = 0.15$) for the 1984-1995 period and ($t = -2.16$, $P = 0.031$) with d.f. = 365 for the 2003-2010 period of trends in N:Extended Data Figure 6. The emphasis in the paper is on the fact that the one trend is significant and the other is not. However the “difference between *significant* and *non-significant* is not itself statistically significant” (Gelman and Stern 2006). The appropriate analysis to do in this context is a test on the difference in effects between the two periods, in general, by comparing a model of main effects and interaction effects between period and imid with a model without the interaction. As the measurements are from the same plot, plots should be added to the model as well. The equivalent simpler analysis is to calculate the difference of the trend estimates between periods for each plot and analyse these differences with the same model as the original data. A significant relation with imid then shows the difference in effect in the two periods.

To the reader and me, only the t-values are available. As a crude proxy of the recommended analysis I reason as follows. As the number of degrees of freedom is large, the distribution of the t-value is close to the standard normal distribution. The difference in t-values is $2.16 - 1.42 = 0.73$ with, assuming independence of the two t-values, a variance of two. This is thus a small difference that is not statistical significant. However, as the data are from the same plots, the t-values are positively correlated giving a smaller variance. With correlation $r > 0$, the variance is $2(1-r)$. For quite high r , namely $r > 1 - 0.73^2/8 = 0.93$, the difference of 0.73 is significant at the 5% level.

This part of the statistical analysis can easily be improved upon with access to the original data analyzed. Such data were not available for this evaluation.

Accounting for other factors (landscape variables)

The part of the paper on accounting for landscape variables is strong from a statistical point of view. The significant of the imid effect is judged in the context of the full multiple mixed effect model for trend using imid and 8 landscape variables as predictors and compared to the effects of the landscape variables (N:Figure 2 and N:Extended Data Table 2). According to N:Sp11, the F statistics in N:Figure 2a are based on single term deletions from the full model (N:Extended Data Table 2). In N:Figure 2a this is described as the marginal variance ratio, which I found confusing; to me it a conditional variance ratio as it is conditional on the other predictors in the model.

A noteworthy result of this analysis is that only imid and Bulb area have significant effects in the model. But note that the effect of Bulb area is *positive* and only marginally significant; it is inferred to increase the local trend. If imid is high in a region with high Bulb area, their joint effect is thus less than predicted by imid alone.

The negative effect of imid in the full model is less in absolute value (namely -0.0294 s.e.m. 0.0077) than in the analyses with imid alone (namely -0.0374 s.e.m. 0.0066).

The analyses on accounting for landscape variables also show that the pseudo-replication problem should not be overemphasized; effects of the eight tested variables (except one) are not significant despite the pseudo-replication mentioned in this section and the section “Issues when inferring from spatial patterns”.

Analyses per species

The analyses per species do not suffer from the pseudo-replication in the sense of the subsection “The joint statistical analysis on all 15 species (LMER)”.

The analyses per species are least-squares regression “weighted by the mean number of territories of plots ... in order to avoid a high influence of demographic stochasticity of small populations” (N:Sp9). In the functions `lm` for regression in the software environment R, which was presumably used for this, the weight total does not influence the outcome of the model including standard errors, so that using the numbers as the weights does not incur any (additional) pseudo-replication. (The function `lmer` for mixed model regression works the same in this respect as the function `lm` for regression). It is nevertheless wiser to write “weighted proportional to the mean number” as in other software a weight of 2 may mean the same as putting the data row twice in the data set.

All but one of the 15 tested species showed a negative slope with respect to imid, and 6 out of 15 are significant at the 5% significance level after Bonferroni correction for multiple testing (Np1).

The kind of possible pseudo-replication in these analyses is the spatial aspect of the data. See the section “Issues when inferring from spatial patterns”.

4. Re-analysis of the binned data

4.1 The binned data

Only a limited amount of information and data is available for this evaluation. The discretized data presented in N:Figure 1b is available from the Nature web site (http://www.nature.com/nature/journal/v511/n7509/source_data/nature13531-f1.csv)

The file also contains the species name attached to each point in the figure and thus allows to check the analysis albeit on binned data. Binning is not necessarily a bad thing as the interpolated imid data contain error anyway. All geographic information was not available, nor information on co-occurrence of species at the same plot.

Figure 1 shows a standard boxplot of the data.

It should be noted that there is one data point not shown in N:Figure 1b, but visible in the boxplot shown in Figure 1: a trend of -0.9775 for the *Mistle thrush* in a single plot (n=1) in the highest dosage (ln-conc = 6.25, conc = 518). Removing this point [or resetting to about the next lowest trend in this concentration class (-0.2)] turns out

not alter the main conclusions of the re-analyses carried out in this subsection (Table 1).

The concerns about pseudo-replication as expressed in section 3 do not apply as much to these analyses as to the original analyses. The number of units in these analyses is far less than in the original analysis (119 against 1459) and even less than the number of monitoring plots (119 against 354). The problem that trends of species in the same plot may be dependent (given the covariates in the model) is not solved in the binned data, but its putative effect is likely to be small. The problem of spatial autocorrelation remains implicit in the analyses of the binned data, but its effect is likely to be small.

The concerns about pseudo-replication as expressed in section 3 do not apply to the analyses that I performed on just nine points, which are weighted median trends in the nine imid classes.

I carried out the LMER analyses in `lme4` which is the successor to the `nlme` library used in the paper and by the same senior author (Bates). All analyses were carried out in R with mention of package name if not `lme4`.




Figure 1. Boxplot of the binned data (unweighted).

4.2 Analyses of the binned data and the median.

Table 1 summarizes the results of the analyses. In Table 1, Model N is the model of the paper as reported in N:Figure 1b.

Model A (Table 1) mimics the analysis in N:Figure 1b as much as possible using the available binned data. It uses 119 observations which are means trends within the bins per species. The weight (W) used is proportional to the number of monitoring plots of the particular species falling a particular imid bin. The conclusions to be drawn from this analysis are identical to those of the original analysis.

The model without imid but with random intercepts for species has an overall mean close to 0, namely -0.00825 (0.00945). The mean bird trend averaged over all birds thus close to 0.

Related models with correlated or uncorrelated random intercept and random slope yield about the same results as model A. The standard deviation of the random slopes is very small compared to the absolute main effect (the slope of -0.034). There is thus no evidence of a difference among species in terms of the effect of imid. A same conclusion is reached when using multiple regression using the model species + species.x (with x is log imid).

Model B (Table 1) ignores the number of plots that is used as weight in Model A and model C (Table 1) removes the outlier in the highest imid class. The outlier has an influence on the unweighted analysis B. It has little influence on the weighted analysis (result not shown). The conclusions to be drawn from these analysis are identical to those of the original analysis.

Model D (Table 1) is suggested by N:Figure 1a, where the contour lines appear to focus on regions with concentrations above and below 50ng/l. This model addresses the issue whether the linear model applies across the whole range of log-imid concentrations. The model consists of two straight lines that connect at 50ng/l. It included random intercepts for species. The slope with respect to imid is very negative and significant in the region below 50ng/l, and less negative and non-significant in the region above 50ng/l. However, as difference between *significant* and *non-significant* may not itself be statistically significant, I tested for the difference and the difference is not statistically significant. The model with a single slope describes the data sufficiently well.

Removing the highest imid class or highest two imid classes strengthens in this data the negative relationship with imid (see Appendix). But the lowest imid values also have a big influence on a regression and these may be due to edge effects in the interpolation method and/or appear in isolated regions that might be peculiar. Model E is an analyses without the lowest imid concentration class (<7.2ng/l) and is very similar to the results of Model A. Model F removes all regions with less than 12ng/l (an idea about the regions involved can be obtained by inspecting the 10ng/l contour lines in N:Fig1a). The result shows a highly significant effect; the regression coefficient is -0.027.

Models G-J are fits on weighted medians trends per concentration class. These are fits on only nine points. The models differ in weighting of the points. The P-values are two-sided. As one does not expect the trend to increase with imid, a better test would be one-sided. For this, the reported P-value must be halved, yielding a significant negative slope at $P < 0.05$, except for model F which is weighted by number of species (which gives low weight to the lowest concentration class).

I included three smooth fits to the binned data (n=119) using the packages `mgcv` and `scam` and using Bayesian P-splines fitted with the INLA and the (not yet public) `BayesPspline` package, but these smooth fits do not fit statistically much better than the straight line model of the paper. Figure 2 shows the fit by `scam` which allows monotonic decreasing fits. The model fit is thus either a horizontal line or, in our case, a decreasing curve (Figure 2). The confidence bands shown in Figure 2 are likely too wide as the model ignored the random intercepts for species, which increases the residual variance of the model.

Table 1. Results of lmer analyses on the binned data.

L ¹	Short description	N ²	W ³	Intercept ⁴	Slope ⁵	s.e.(slope) ⁶	P ⁷
N	N:Fig.1b	1459	1	0.110	-0.0374	0.0066	<0.0001
A	Weighted binned data	119	<i>n</i>	0.100	-0.0342	0.0076	<0.0001
B	Unweighted binned data	119	1	0.124	-0.0421	0.0107	0.00014
C	Unweighted without the outlier	118	1	0.093	-0.0319	0.0094	0.00091
D	Weighted, change point at 50ng/l?	119	<i>n</i>	<50ng/l >50ng/l difference	-0.0398 -0.0243 0.0154	0.0122 0.0190 0.0270	0.0011 0.02 0.57
E	Weighted without 7.4ng/l	112	<i>n</i>	0.097	-0.0335	0.0079	0.00005
F	Weighted without <12.2 ng/l	97	<i>n</i>	0.070	-0.0268	0.0100	0.0092
	Analyses (G-J)		on	weighted	medians		
G	Unweighted	9	1	0.0719	-0.0233	0.0118	0.089
H	Weighted by number species	9	<i>s</i>	0.0513	-0.0191	0.0116	0.14
I	Weighted by number plots	9	<i>n</i>	0.05568	-0.0215	0.0099	0.065
J	Weighted by inverse variance	9	1/var(y)	0.0830	-0.0309	0.0103	0.02

¹Model labels : N- original analysis, A: mimic of N using the binned data, B: as A but not using weights for the data points, C: as B, but without the outlier in the highest imid class, D: a change point model with a slope for log-imid <50ng/l and a slope for log-imid >50ng/l, E: analysis as A but without log-imid <7.4ng/l, D: analysis as A but without log-imid <12.2 ng/l, G-J: simple linear regressions on the weighted medians per imid bin of N:Fig1b on log imid.

²N: number of observations. ³W: weight of each observation (*n* = number of monitoring plots or 1).

⁴intercept: estimate of the intercept. ⁵slope: estimate of the slope with respect to log imid. ⁶s.e. (slope): standard error of the slope estimate.

⁷P: P-value of statistical test of the null hypothesis that the slope is equal to zero (no effect of log imid, or, in Model D, a zero difference between slopes).

Decreasing fit with 95% confidence intervals




Figure 2. Unweighted smooth fit using package `scam`. The dotted line gives the expected trend at the lowest concentration class and can be considered as a reference line (see subsection Interpretation of the joint analysis). The model did not include random intercepts for species. Consequently the confidence bands may be too wide.

4.3 Conclusions for the analyses of the binned data

The conclusion from the analyses on the binned data is that the statistical relationship between the bird trends and imid is robust against a number of appropriate changes in the data and analysis.

Disclaimer

This report is made without access to the original data and on short notice.

Conflict of interest

The author is co-author on a paper in [5.1.2.e](#) with one of the authors of the Nature paper [5.1.2.e](#). The author does contract research with industry, but not currently or in the recent past with Bayer. These points are mentioned for transparency purposes.

5. References

- Bates, D., M. Maechler, and B. Bolker. 2014. lme4: Linear mixed-effects models using S4 classes. R package version 1.1-7, <http://CRAN.R-project.org/package=lme4>.
- Dormann, C. F., J. M. McPherson, M. B. Araújo, R. Bivand, J. Bolliger, G. Carl, R. G. Davies, A. Hirzel, W. Jetz, W. D. Kissling, I. Kühn, R. Ohlemüller, P. R. Peres-Neto, B. Reineking, B. Schröder, F. M. Schurr, and R. Wilson. 2007. Methods to account for spatial autocorrelation in the analysis of species distributional data: a review. *Ecography* **30**:609-628.
- Gelman, A. and H. Stern. 2006. The Difference Between “Significant” and “Not Significant” is not Itself Statistically Significant. *The American Statistician* **60**:328-331.
- Hallmann, C. A., R. P. B. Foppen, C. A. M. van Turnhout, H. de Kroon, and E. Jongejans. 2014. Declines in insectivorous birds are associated with high neonicotinoid concentrations. *Nature* **511**:341-343.
- Hurlbert, S. H. 1984. Pseudoreplication and the design of ecological field experiments. *Ecological Monographs* **54**:187-211.
- Jamil, T., W. A. Ozinga, M. Kleyer, and C. J. F. ter Braak. 2013. Selecting traits that explain species–environment relationships: a generalized linear mixed model approach. *Journal of Vegetation Science* **24**:988-1000.
- Jongman, R. H. G., C. J. F. ter Braak, and O. F. R. van Tongeren. 1995. Data analysis in community and landscape ecology. Cambridge University Press, Cambridge.
- Muff, S., A. Riebler, L. Held, H. Rue, and P. Saner. 2014. Bayesian analysis of measurement error models using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*:n/a-n/a.
- Rubin, D. B. 1996. Multiple Imputation After 18+ Years. *Journal of the American Statistical Association* **91**:473-489.
- van Buuren, S. 2012. Flexible Imputation of Missing Data. Chapman & Hall/CRC, Boca Raton, FL.
- Vijver, M. G. and P. J. van den Brink. 2014. Macro-Invertebrate Decline in Surface Water Polluted with Imidacloprid: A Rebuttal and Some New Analyses. *Plos one* **9**:e89837.

6. Appendix R code with results

6.1 Analysis of the binned data (models A-F)

Fig1b_Binned_data.r

5.1.2.e

Thu Aug 14 16:49:40 2014

```
rm(list=ls(all=TRUE))
setwd("D:/Data/Aa/_projects/Ctgb")
nature13531.f1 <- read.csv("D:/Data/Aa/_projects/Ctgb/nature13531-f
1.csv")
Data = nature13531.f1
#View(Data)
varnames.ori = names(nature13531.f1)
Data1 = Data
names(Data1)[2:5]= c("y", "x", "species", "Conc.class")
Data = data.frame(Data,Data1)

names(Data)

## [1] "X"                "y_point"          "x_point"
## [4] "Spec."            "Concentration_class" "n"
## [7] "X.1"              "y"                "x"
## [10] "species"          "Conc.class"       "n.1"

# n = number of plot combinations
# Each point represents the average intrinsic rate of increase of a
species
# over all plots in the same concentration class, whereas
# the size of the point is scaled proportionally to the number of sp
ecies plot combinations on which the calculated mean is based. Binni
ng into classes was performed to reduce scatter noise and aid in vis
ual interpretation. Actual analysis, and the depicted regression lin
e, was performed on raw data (n = 1,459). The regression line is giv
en by  $0.1110 - 0.0374$  (s.e.m. = 0.0066)  $\times \log[\text{imidacloprid}]$  ( $P < 0.0001$ ).
# Dashed lines delineate the 95% confidence interval.
sum(Data$n) # 1459 OK

## [1] 1459

rem1.1 = FALSE

attach(Data)
summary(Data)

##          X          y_point          x_point
Spec.
## Min.    : 2.0   Min.    :-0.9775   Min.    :1.75   Common whitethr
oat: 9
```

```


## 1st Qu.: 38.5  1st Qu.: -0.0808  1st Qu.: 2.75  Marsh Warble
r      : 9
## Median : 68.0  Median : -0.0215  Median : 3.75  Meadow pipi
t      : 9
## Mean   : 69.8  Mean   : -0.0372  Mean   : 3.86  Reed warble
r      : 9
## 3rd Qu.:101.5  3rd Qu.: 0.0447  3rd Qu.: 4.75  Stonecha
t      : 9
## Max.   :135.0  Max.   : 0.3703  Max.   : 6.25  Willow Warble
r      : 9
##
##                                     (Othe
r)                                     :65
## Concentration_class      n      X.1      y
## (2,2.5]:15      Min.   : 1.0  Min.   : 2.0  Min.   : -0.97
75
## (2.5,3]:15      1st Qu.: 2.0  1st Qu.: 38.5  1st Qu.: -0.08
08
## (3,3.5]:15      Median : 7.0  Median : 68.0  Median : -0.02
15
## (3.5,4]:15      Mean   :12.3  Mean   : 69.8  Mean   : -0.03
72
## (4.5,5]:15      3rd Qu.:16.5  3rd Qu.:101.5  3rd Qu.: 0.04
47
## (5,5.5]:13      Max.   :65.0  Max.   :135.0  Max.   : 0.37
03
## (Other):31
##      x      species      Conc.class      n.1
## Min.   :1.75  Common whitethroat: 9  (2,2.5]:15  Min.   : 1.0
## 1st Qu.:2.75  Marsh Warbler      : 9  (2.5,3]:15  1st Qu.: 2.0
## Median :3.75  Meadow pipit       : 9  (3,3.5]:15  Median : 7.0
## Mean   :3.86  Reed warbler       : 9  (3.5,4]:15  Mean   :12.3
## 3rd Qu.:4.75  Stonechat          : 9  (4.5,5]:15  3rd Qu.:16.5
## Max.   :6.25  Willow Warbler     : 9  (5,5.5]:13  Max.   :65.0
##                                     (Other) :65  (Other):31

specn = species
levels(specn) = 1:nlevels(species)


my.ylab = "Mean species trend"
my.xlab = "ln(Imid. Conc)"
ylim1 = c(-1,0.4) # as in paper
ylim2 = c(-0.25,0.25)
ylim3 = c(-0.6,0.4)
check.outlier = TRUE

plot(x,y,pch=".", ylab = my.ylab, xlab= my.xlab)
symbols(x,y, circles = n/700,inches= FALSE, add= TRUE)
abline(0,0, lty = "dotted")

```

```
# default boxplots
boxplot(y~Data$Conc.class,ylab = my.ylab, xlab= "lmid log conc. class")
abline(0,0, lty = "dotted")
```



```
levels.x = unique(x_point);
levels.conc = round(exp(unique(x)),1);
```

```

Conc.int= levels(Conc.class)
class = 1:nlevels(Conc.class)
data.frame(class,levels.x,levels.conc, Conc.int)

##   class levels.x levels.conc Conc.int
## 1     1     1.75         5.8 (1.5,2]
## 2     2     2.25         9.5 (2,2.5]
## 3     3     2.75        15.6 (2.5,3]
## 4     4     3.25        25.8 (3,3.5]
## 5     5     3.75        42.5 (3.5,4]
## 6     6     4.25        70.1 (4,4.5]
## 7     7     4.75       115.6 (4.5,5]
## 8     8     5.25       190.6 (5,5.5]
## 9     9     6.25       518.0 (5.5,7]

conc = exp(x)
bins = levels(Conc.class)
table(Conc.class)

## Conc.class
## (1.5,2] (2,2.5] (2.5,3] (3,3.5] (3.5,4] (4,4.5] (4.5,5] (5,5.5]
## (5.5,7]
##      7      15      15      15      15      11      15      13
## 13

library(lme4)

## Loading required package: Matrix
## Loading required package: Rcpp

#####
#
#####
#####
# lmer analyses on 119 observations
#####

Data$x = Data$x_point
Data$w = Data$n/sum(Data$n)*nrow(Data) # normalization not need, same result with Data$w = Data$n
mod0 = lmer(y~ (1| species), data = Data, REML = reml.1, weights = w)
summary(mod0)

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ (1 | species)
##   Data: Data
## Weights: w
##
##      AIC      BIC    logLik deviance df.resid
## -172.0   -163.7     89.0   -178.0     116
##
## Scaled residuals:
##   Min      1Q  Median      3Q      Max
## -3.408 -0.736 -0.078  0.428  2.445
##

```



```

## Random effects:
## Groups Name Variance Std.Dev.
## species (Intercept) 0.000421 0.0205
## Residual 0.006542 0.0809
## Number of obs: 119, groups: species, 15
##
## Fixed effects:
## Estimate Std. Error t value
## (Intercept) -0.00825 0.00945 -0.87

#####
# model A
mod1 = lmer(y~ x + (1 | species), data = Data, REML = reml.1, weight
s = w)
summary(mod1)

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ x + (1 | species)
## Data: Data
## Weights: w
##
## AIC BIC logLik deviance df.resid
## -188.2 -177.1 98.1 -196.2 115
##
## Scaled residuals:
## Min 1Q Median 3Q Max
## -3.308 -0.622 0.006 0.607 2.472
##
## Random effects:
## Groups Name Variance Std.Dev.
## species (Intercept) 0.000586 0.0242
## Residual 0.005485 0.0741
## Number of obs: 119, groups: species, 15
##
## Fixed effects:
## Estimate Std. Error t value
## (Intercept) 0.10024 0.02605 3.85
## x -0.03429 0.00767 -4.47
##
## Correlation of Fixed Effects:
## (Intr)
## x -0.929

#####
midx = weighted.mean(Data$x_point,w = Data$n) # 3.86 unweighted, 3.1
4 weighted
midx

## [1] 3.146

Data$x = Data$x_point -midx # explicit centering in case lmer does
not do so.
mod2 = lmer(y~ x + (1|species) + (0+x| species), data = Data, REML =
reml.1, weights = w)
summary(mod2)

```

```

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ x + (1 | species) + (0 + x | species)
## Data: Data
## Weights: w
##
##      AIC      BIC   logLik deviance df.resid
## -186.2 -172.3   98.1  -196.2    114
##
## Scaled residuals:
##   Min     1Q  Median     3Q      Max
## -3.308 -0.622  0.006  0.607  2.472
##
## Random effects:
## Groups      Name          Variance Std.Dev.
## species    (Intercept) 0.000586 0.0242
## species.1 x              0.000000 0.0000
## Residual                0.005485 0.0741
## Number of obs: 119, groups: species, 15
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept) -0.00765    0.00962  -0.79
## x            -0.03429    0.00767  -4.47
##
## Correlation of Fixed Effects:
## (Intr)
## x -0.007

mod3 = lmer(y~ x + (1 + x| species), data = Data, REML = reml.1, wei
ghts = w)
summary(mod3)

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ x + (1 + x | species)
## Data: Data
## Weights: w
##
##      AIC      BIC   logLik deviance df.resid
## -184.3 -167.6   98.1  -196.3    113
##
## Scaled residuals:
##   Min     1Q  Median     3Q      Max
## -3.322 -0.610  0.010  0.615  2.530
##
## Random effects:
## Groups      Name          Variance Std.Dev. Corr
## species    (Intercept) 5.74e-04 0.02395
##           x              9.24e-06 0.00304 -1.00
## Residual                5.48e-03 0.07405
## Number of obs: 119, groups: species, 15
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept) -0.00714    0.00958  -0.75

```



```

## x          -0.03447    0.00772   -4.46
##
## Correlation of Fixed Effects:
##   (Intr)
## x -0.084

anova(mod1,mod2,mod3,mod0)

## Data: Data
## Models:
## mod0: y ~ (1 | species)
## mod1: y ~ x + (1 | species)
## mod2: y ~ x + (1 | species) + (0 + x | species)
## mod3: y ~ x + (1 + x | species)
##      Df  AIC  BIC logLik deviance Chisq Chi Df Pr(>Chisq)
## mod0  3 -172 -164  89.0     -178      18.2    1  2e-05 ***
## mod1  4 -188 -177  98.1     -196    0.0    1  1.00
## mod2  5 -186 -172  98.1     -196    0.1    1  0.76
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

anova(mod0,mod1)

## Data: Data
## Models:
## mod0: y ~ (1 | species)
## mod1: y ~ x + (1 | species)
##      Df  AIC  BIC logLik deviance Chisq Chi Df Pr(>Chisq)
## mod0  3 -172 -164  89.0     -178      18.2    1  2e-05 ***
## mod1  4 -188 -177  98.1     -196    0.1    1  0.76
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

anova(mod0,mod2)

## Data: Data
## Models:
## mod0: y ~ (1 | species)
## mod2: y ~ x + (1 | species) + (0 + x | species)
##      Df  AIC  BIC logLik deviance Chisq Chi Df Pr(>Chisq)
## mod0  3 -172 -164  89.0     -178      18.2    2  0.00011 ***
## mod2  5 -186 -172  98.1     -196    0.1    1  0.76
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

anova(mod0,mod3)

## Data: Data
## Models:
## mod0: y ~ (1 | species)
## mod3: y ~ x + (1 + x | species)
##      Df  AIC  BIC logLik deviance Chisq Chi Df Pr(>Chisq)
## mod0  3 -172 -164  89.0     -178      18.3    3  0.00039 ***

```

```

## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

# reinstall uncentred x
Data$x = Data$x_point

lmod0 = lm(y~species, data = Data, weights = w)
lmod1 = lm(y~species + x, data = Data, weights = w)
lmod2 = lm(y~species + species:x, data = Data, weights = w)
# summary(Lmod1)
# summary(Lmod2)
anova(lmod0,lmod1,lmod2)

## Analysis of Variance Table
##
## Model 1: y ~ species
## Model 2: y ~ species + x
## Model 3: y ~ species + species:x
##   Res.Df  RSS Df Sum of Sq    F Pr(>F)
## 1     104 0.696
## 2     103 0.582  1    0.114 21.18 1.4e-05 ***
## 3      89 0.481 14    0.101  1.33  0.21
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

# x has an effect but the slope does not differ significant between
species

if (check.outlier){
# Look at the effect of the outlier in model A
ii =which.min(y)
iii= which.min(y[-ii])
c(ii,y[ii], x[ii],n[ii],species[ii]) # outlier
species[ii]
c(iii,y[iii], x[iii],n[iii],species[iii])
species[iii]
yii = Data$y[ii]; yii
Data118 = Data[-ii,] # remove the outlier
mod0 = lmer(y~ (1| species), data = Data118, REML = reml.1, weights
= w)
summary(mod0)
mod1 = lmer(y~ x + (1 | species), data = Data118, REML = reml.1, wei
ghts = w)
print(summary(mod1))
print(anova(mod0,mod1))
# the effect of the outlier is very minor in the weighted analysis
}

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ x + (1 | species)
##   Data: Data118
## Weights: w
##

```

```

##      AIC      BIC  logLik deviance df.resid
##    -199.0   -187.9   103.5   -207.0     114
##
## Scaled residuals:
##      Min      1Q  Median      3Q      Max
## -2.9295 -0.6830  0.0134  0.6251  2.5381
##
## Random effects:
##  Groups   Name                Variance Std.Dev.
## species (Intercept) 0.000618 0.0249
## Residual                0.004971 0.0705
## Number of obs: 118, groups: species, 15
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept)  0.09405    0.02501    3.76
## x            -0.03201    0.00734   -4.36
##
## Correlation of Fixed Effects:
##  (Intr)
## x -0.925
## Data: Data118
## Models:
## mod0: y ~ (1 | species)
## mod1: y ~ x + (1 | species)
##      Df  AIC  BIC logLik deviance Chisq Chi Df Pr(>Chisq)
## mod0  3 -184 -175   94.8   -190
## mod1  4 -199 -188  103.5   -207  17.4    1  3.1e-05 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

#####
# model B
# unweighted analysis
Data$x = Data$x_point
mod0 = lmer(y~ (1| species), data = Data, REML = reml.1)
summary(mod0)

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ (1 | species)
## Data: Data
##
##      AIC      BIC  logLik deviance df.resid
##    -83.2   -74.9   44.6   -89.2     116
##
## Scaled residuals:
##      Min      1Q  Median      3Q      Max
## -5.618 -0.257  0.104  0.518  2.421
##
## Random effects:
##  Groups   Name                Variance Std.Dev.
## species (Intercept) 0.000536 0.0231
## Residual                0.027157 0.1648
## Number of obs: 119, groups: species, 15

```



```

##
## Fixed effects:
##           Estimate Std. Error t value
## (Intercept) -0.0379    0.0163  -2.33

mod1 = lmer(y~ x + (1 | species), data = Data, REML = reml.1)
summary(mod1)

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ x + (1 | species)
##   Data: Data
##
##      AIC      BIC    logLik deviance df.resid
##   -95.7   -84.6     51.8  -103.7     115
##
## Scaled residuals:
##   Min      1Q  Median      3Q      Max
## -5.288 -0.400  0.006  0.435  2.789
##
## Random effects:
##   Groups   Name      Variance Std.Dev.
##   species (Intercept) 0.000913 0.0302
##   Residual              0.023686 0.1539
## Number of obs: 119, groups:  species, 15
##
## Fixed effects:
##           Estimate Std. Error t value
## (Intercept)  0.1243    0.0444   2.80
## x            -0.0421    0.0107  -3.93
##
## Correlation of Fixed Effects:
##   (Intr)
## x -0.931

anova(mod0,mod1)

## Data: Data
## Models:
## mod0: y ~ (1 | species)
## mod1: y ~ x + (1 | species)
##      Df   AIC   BIC logLik deviance Chisq Chi Df Pr(>Chisq)
## mod0  3 -83.2 -74.9  44.6   -89.2
## mod1  4 -95.7 -84.6  51.8  -103.7  14.5    1  0.00014 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

if (check.outlier){
#####
# model C, unweighted without outlier
mod0 = lmer(y~ (1| species), data = Data118, REML = reml.1)
summary(mod0)
mod1 = lmer(y~ x + (1 | species), data = Data118, REML = reml.1)
print(summary(mod1))
print(anova(mod0,mod1))
}

```

```

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ x + (1 | species)
## Data: Data118
##
## AIC      BIC    logLik deviance df.resid
## -128.1   -117.0    68.0  -136.1    114
##
## Scaled residuals:
##   Min      1Q  Median      3Q      Max
## -3.591 -0.447  0.002  0.466  2.966
##
## Random effects:
## Groups   Name              Variance Std.Dev.
## species (Intercept) 0.00087  0.0295
## Residual              0.01773  0.1332
## Number of obs: 118, groups: species, 15
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept)  0.0929     0.0389    2.39
## x            -0.0320     0.0094   -3.40
##
## Correlation of Fixed Effects:
## (Intr)
## x -0.928
## Data: Data118
## Models:
## mod0: y ~ (1 | species)
## mod1: y ~ x + (1 | species)
##      Df  AIC  BIC loglik deviance Chisq Chi Df Pr(>Chisq)
## mod0  3 -119 -111  62.5   -125
## mod1  4 -128 -117  68.0   -136   11    1  0.00091 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

#####
# Change point at 50ng/L?
#####

x50min = Data$x - log(50)
x50min[x50min >0 ] =0 # all greater values of x set to 0
Data$addx50min = x50min
x50plus = Data$x - log(50)
x50plus[x50plus < 0 ] =0
Data$addx50plus = x50plus
Data$w = Data$n/sum(Data$n)*nrow(Data) # normalization not need, sam
e result with Data$w = Data$n

#####
# model D
mod1 = lmer(y~ x + (1 | species), data = Data, REML = reml.1, weight
s = w)
#summary(mod1)
mod2 = lmer(y~ x + addx50plus + (1 | species), data = Data, REML =

```

```

reml.1, weights = w)
summary(mod2)

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ x + addx50plus + (1 | species)
## Data: Data
## Weights: w
##
##      AIC      BIC    logLik deviance df.resid
## -186.5 -172.6     98.3  -196.5     114
##
## Scaled residuals:
##   Min     1Q Median     3Q      Max
## -3.380 -0.636  0.025  0.582  2.418
##
## Random effects:
## Groups   Name              Variance Std.Dev.
## species (Intercept) 0.000592 0.0243
## Residual              0.005466 0.0739
## Number of obs: 119, groups: species, 15
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept)  0.1154     0.0371    3.11
## x            -0.0398     0.0122   -3.25
## addx50plus   0.0154     0.0270    0.57
##
## Correlation of Fixed Effects:
##              (Intr) x
## x            -0.964
## addx50plus   0.713 -0.780

# the coefficient for x in mod2 is the slope for <50ng/L
(Pmin50 =2*pnorm(-3.25))

## [1] 0.001154

mod3 = lmer(y~ x + addx50min + (1 | species), data = Data, REML = r
eml.1, weights = w)
summary(mod3)

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ x + addx50min + (1 | species)
## Data: Data
## Weights: w
##
##      AIC      BIC    logLik deviance df.resid
## -186.5 -172.6     98.3  -196.5     114
##
## Scaled residuals:
##   Min     1Q Median     3Q      Max
## -3.380 -0.636  0.025  0.582  2.418
##
## Random effects:
## Groups   Name              Variance Std.Dev.

```



```

## species (Intercept) 0.000592 0.0243
## Residual            0.005466 0.0739
## Number of obs: 119, groups: species, 15
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept)  0.0551    0.0832    0.66
## x            -0.0243    0.0190   -1.28
## addx50min    -0.0154    0.0270   -0.57
##
## Correlation of Fixed Effects:
##              (Intr) x
## x            -0.986
## addx50min    0.950 -0.915

# the coefficient for x in mod3 is the slope for >50ng/L
(Pplus50 =2*pnorm(-1.28))

## [1] 0.2005

anova(mod1,mod2)

## Data: Data
## Models:
## mod1: y ~ x + (1 | species)
## mod2: y ~ x + addx50plus + (1 | species)
##      Df  AIC  BIC loglik deviance Chisq Chi Df Pr(>Chisq)
## mod1  4 -188 -177  98.1    -196
## mod2  5 -186 -173  98.3    -196  0.33    1    0.57

# model 2 not statistically better than model 1
# check
anova(mod1,mod3)

## Data: Data
## Models:
## mod1: y ~ x + (1 | species)
## mod3: y ~ x + addx50min + (1 | species)
##      Df  AIC  BIC loglik deviance Chisq Chi Df Pr(>Chisq)
## mod1  4 -188 -177  98.1    -196
## mod3  5 -186 -173  98.3    -196  0.33    1    0.57

# no significant difference in slope

#####
# model E
Data_1 = Data[x> min(x),] # data without the first imid bin
dim(Data_1) # 112 by 15

## [1] 112 15

summary(Data_1$x)

##      Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
##      2.25   2.75   3.75   4.00   4.75   6.25

```

```

mod0 = lmer(y~ (1| species), data = Data_1, REML = reml.1, weights
= w)
summary(mod0)

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ (1 | species)
## Data: Data_1
## Weights: w
##
## AIC      BIC    logLik deviance df.resid
## -170.6   -162.5    88.3  -176.6    109
##
## Scaled residuals:
##  Min      1Q  Median      3Q      Max
## -3.348 -0.746 -0.149  0.411  2.407
##
## Random effects:
## Groups Name          Variance Std.Dev.
## species (Intercept) 0.000407 0.0202
## Residual              0.006777 0.0823
## Number of obs: 112, groups: species, 15
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept) -0.00888    0.00951  -0.93

mod1 = lmer(y~ x + (1 | species), data = Data_1, REML = reml.1, weig
hts = w)
print(summary(mod1))

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ x + (1 | species)
## Data: Data_1
## Weights: w
##
## AIC      BIC    logLik deviance df.resid
## -185.1   -174.2    96.5  -193.1    108
##
## Scaled residuals:
##  Min      1Q  Median      3Q      Max
## -3.252 -0.631 -0.002  0.573  2.417
##
## Random effects:
## Groups Name          Variance Std.Dev.
## species (Intercept) 0.000574 0.0240
## Residual              0.005709 0.0756
## Number of obs: 112, groups: species, 15
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept)  0.09734    0.02672   3.64
## x            -0.03349    0.00788  -4.25
##
## Correlation of Fixed Effects:

```

```

## (Intr)
## x -0.932

print(anova(mod0,mod1))

## Data: Data_1
## Models:
## mod0: y ~ (1 | species)
## mod1: y ~ x + (1 | species)
##      Df  AIC  BIC logLik deviance Chisq Chi Df Pr(>Chisq)
## mod0  3 -171 -162  88.3    -177
## mod1  4 -185 -174  96.5    -193  16.5    1  4.9e-05 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

# model F
Data_2 = Data[x > 2.25,] # data without the first two imid bins
dim(Data_2) # 97 by 15

## [1] 97 15

summary(Data_2$x)

##      Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
##      2.75   3.25   4.25   4.27   5.25   6.25

mod0 = lmer(y~ (1| species), data = Data_2, REML = reml.1, weights
= w)
summary(mod0)

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ (1 | species)
## Data: Data_2
## Weights: w
##
##      AIC      BIC    logLik deviance df.resid
## -138.1 -130.4    72.0  -144.1     94
##
## Scaled residuals:
##      Min      1Q  Median      3Q      Max
## -3.375 -0.719 -0.138  0.343  2.084
##
## Random effects:
## Groups Name Variance Std.Dev.
## species (Intercept) 0.000218 0.0148
## Residual 0.006489 0.0806
## Number of obs: 97, groups: species, 15
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept) -0.0234    0.0096  -2.44

mod1 = lmer(y~ x + (1 | species), data = Data_2, REML = reml.1, weig
hts = w)
print(summary(mod1))

```



```

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ x + (1 | species)
## Data: Data_2
## Weights: w
##
##      AIC      BIC    logLik deviance df.resid
##   -142.9   -132.6     75.4   -150.9     93
##
## Scaled residuals:
##   Min      1Q  Median      3Q      Max
## -3.240 -0.613 -0.046  0.580  2.394
##
## Random effects:
## Groups Name          Variance Std.Dev.
## species (Intercept) 0.000318 0.0178
## Residual              0.005969 0.0773
## Number of obs: 97, groups: species, 15
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept)  0.0694    0.0362    1.92
## x            -0.0268    0.0100   -2.66
##
## Correlation of Fixed Effects:
##   (Intr)
## x -0.963

print(anova(mod0,mod1))

## Data: Data_2
## Models:
## mod0: y ~ (1 | species)
## mod1: y ~ x + (1 | species)
##      Df  AIC  BIC logLik deviance Chisq Chi Df Pr(>Chisq)
## mod0  3 -138 -130  72.0   -144
## mod1  4 -143 -133  75.4   -151  6.79    1  0.0092 **
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

#
#####

# model XX
Data_89 = Data[x < 5,] # data without the last two imid bins
dim(Data_89) # 97 by 15

## [1] 93 15

summary(Data_89$x)

##   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
##   1.75  2.75   3.25   3.34  4.25   4.75

mod0 = lmer(y~ (1| species), data = Data_89, REML = reml.1, weights
= w)
summary(mod0)

```

```

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ (1 | species)
## Data: Data_89
## Weights: w
##
##      AIC      BIC   logLik deviance df.resid
## -160.2   -152.6    83.1   -166.2     90
##
## Scaled residuals:
##   Min     1Q  Median     3Q    Max
## -2.834 -0.680 -0.107  0.469  2.401
##
## Random effects:
## Groups   Name      Variance Std.Dev.
## species (Intercept) 0.000438 0.0209
## Residual              0.006539 0.0809
## Number of obs: 93, groups: species, 15
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept) -0.00602    0.00967   -0.62

mod1 = lmer(y~ x + (1 | species), data = Data_89, REML = reml.1, wei
ghts = w)
print(summary(mod1))

## Linear mixed model fit by maximum likelihood ['lmerMod']
## Formula: y ~ x + (1 | species)
## Data: Data_89
## Weights: w
##
##      AIC      BIC   logLik deviance df.resid
## -181.5   -171.4    94.8   -189.5     89
##
## Scaled residuals:
##   Min     1Q  Median     3Q    Max
## -2.7988 -0.7034 -0.0615  0.5503  2.3063
##
## Random effects:
## Groups   Name      Variance Std.Dev.
## species (Intercept) 0.000629 0.0251
## Residual              0.004892 0.0699
## Number of obs: 93, groups: species, 15
##
## Fixed effects:
##              Estimate Std. Error t value
## (Intercept)  0.13948    0.02939   4.75
## x            -0.04769    0.00916  -5.21
##
## Correlation of Fixed Effects:
##   (Intr)
## x -0.945

print(anova(mod0,mod1))

```

```

## Data: Data_89
## Models:
## mod0: y ~ (1 | species)
## mod1: y ~ x + (1 | species)
##      Df  AIC  BIC logLik deviance Chisq Chi Df Pr(>Chisq)
## mod0  3 -160 -153  83.1    -166
## mod1  4 -182 -171  94.8    -190  23.3    1  1.4e-06 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

#

#####
# nonlinear smooth analyses without random species intercept
#####

library(mgcv)

## Loading required package: nlme
##
## Attaching package: 'nlme'
##
## The following object is masked from 'package:lme4':
##
##   lmList
##
## This is mgcv 1.8-2. For overview type 'help("mgcv-package")'.

Data$w = Data$n/sum(Data$n)*nrow(Data) # normalization not need, same result with Data$w = Data$n
gmod = gam(y~s(x, k= 5), data= Data, weights = w )
summary(gmod)

##
## Family: gaussian
## Link function: identity
##
## Formula:
## y ~ s(x, k = 5)
##
## Parametric coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept) -0.0265    0.0102   -2.59   0.011 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Approximate significance of smooth terms:
##      edf Ref.df    F p-value
## s(x) 3.64  3.93 5.85 0.00028 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##


```



```
## R-sq.(adj) = 0.143   Deviance explained = 18%
## GCV = 0.0061993   Scale est. = 0.0059575   n = 119

#plot(gmod,pages=1,residuals=TRUE) ## show partial residuals
my.main = "Smooth fit with 95% confidence intervals"
plot(gmod,pages=1,seWithMean=TRUE, ylab = my.ylab, xlab = my.xlab,
main = my.main)
points(x,y,pch=".", ylab = my.ylab, xlab= my.xlab)
symbols(x,y, circles = n/700,inches= FALSE, add= TRUE)
abline(0,0)
abline(fitted(gmod)[1],0,lty = "dotted")
```

Smooth fit with 95% confidence intervals



```
#####
# decreasing smooth analyses without random species intercept
#####

library(scam)

## This is scam 1.1-7.

#
smod = NULL

# weighed analysis did not run
#smod = scam(y~s(x, k= 8, bs="mpd"), data= Data, weights = w )
smod = scam(y~s(x, k= 8, bs="mpd"), data= Data)


## Loading required package: splines

summary(smod)
```

```
##
## Family: gaussian
## Link function: identity
##
## Formula:
## y ~ s(x, k = 8, bs = "mpd")
##
## Parametric coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept)  0.422      0.596    0.71   0.48
##
## Approximate significance of smooth terms:
##      edf Ref.df    F p-value
## s(x)   3   3.01 8.17 5.3e-05 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## R-sq.(adj) =  0.1545   Deviance explained = 17.6%
## GCV score = 0.024429   Scale est. = 0.023608   n = 119

newdata = data.frame(x= seq(from =min(x), to =max(x), length.out = 1
00), y= NA)
Ps = predict(smod,newdata, se.fit= TRUE)
my.main = "Decreasing fit with 95% confidence intervals"
plot(newdata$x,Ps$fit, type = "l", main = my.main, ylim = ylim2, ylab=
my.ylab, xlab=my.xlab)
lines(newdata$x, Ps$fit+ 2*Ps$se.fit, lty = "dashed" )
lines(newdata$x, Ps$fit- 2*Ps$se.fit, lty = "dashed" )
abline(0,0)
abline(fitted(smod)[1],0,lty = "dotted")
```

Decreasing fit with 95% confidence intervals




```

plot(newdata$x,Ps$fit, type = "l", main = my.main, ylim = ylim3, ylab = my.ylab, xlab=my.xlab)
lines(newdata$x, Ps$fit+ 2*Ps$se.fit, lty = "dashed" )
lines(newdata$x, Ps$fit- 2*Ps$se.fit, lty = "dashed" )
points(x,y,pch=1, ylab = my.ylab, xlab= my.xlab)
abline(0,0)
abline(fitted(smod)[1],0,lty = "dotted")

```

Decreasing fit with 95% confidence intervals



```

#####
# nonlinear smooth analyses with random species intercept
# using Bayesian P-splines
#####

library(INLA)

## Loading required package: sp

library(BayesianPsplines)

a1 = 1; # d = 1
a2 = 0.00001
print(c(a1,a2))

## [1] 1e+00 1e-05

hyper.gamma = list(prec = list(prior = "loggamma", param = c(a1, a
2))) # the default, not very good

mydata = data.frame(y= y,group = species, x= x)
rs = NULL
rs = smooth_inla(mydata, hyperB= hyper.gamma)

```




```

## Note: method with signature 'CsparseMatrix#Matrix#missing#replValue'
## chosen for function '['<-',
## target signature 'dgCMatrix#lgeMatrix#missing#numeric'.
## "Matrix#ldenseMatrix#missing#replValue" would also be valid
## Note: method with signature 'TsparseMatrix#Matrix#missing#replValue'
## chosen for function '['<-',
## target signature 'dgTMatrix#ngCMatrix#missing#numeric'.
## "Matrix#nsparseMatrix#missing#replValue" would also be valid

# intercept
intercept = rs$pred[1,]
k = 1 # first predictor
main.txt = "Bayesian smooth with 95% credible region"
pred = component_plus_residual(rs, k, mplot = TRUE, ylim = ylim3, ylab = my.ylab,
xlab = my.xlab, main = main.txt, pch = 20)
abline(0,0)
abline(pred$pred.grid[1,"mean"],0, lty="dotted")

```

Bayesian smooth with 95% credible region



```

dic1 = rs$model$dic$dic

contr.comp = inla.set.control.compute.default()
contr.comp$dic = TRUE
rs0 = inla(y~x+f(species,model = "iid"), data = list(y= y,x= x, species = species ),control.compute = contr.comp)
summary(rs0)

##
## Call:
## c("inla(formula = y ~ x + f(species, model = \"iid\"), data = list(y = y, \" \" x = x, species = species), control.compute = cont

```

```

r.comp)" )
##
## Time used:
## Pre-processing      Running inla Post-processing      Total
##           0.1470           0.2470           0.1590           0.5531
##
## Fixed effects:
##           mean      sd 0.025quant 0.5quant 0.975quant      mode
kld
## (Intercept) 0.1256 0.0447      0.0378  0.1256      0.2134  0.1256
0
## x           -0.0422 0.0109      -0.0637 -0.0422      -0.0207 -0.0422
0
##
## Random effects:
## Name      Model
## species   IID model
##
## Model hyperparameters:
##           mean      sd      0.025qu
ant
## Precision for the Gaussian observations      40.81      5.31      31.2
1
## Precision for species                        17251.59 17722.09      956.8
9
##           0.5quant 0.975quant mode
## Precision for the Gaussian observations      40.54      52.00      4
0.08
## Precision for species                        11884.86 64259.94      245
1.52
##
## Expected number of effective parameters(std dev): 2.57(0.5836)
## Number of equivalent replicates : 46.31
##
## Deviance Information Criterion: -97.63
## Effective number of parameters: 3.53
##
## Marginal Likelihood: 33.00
## Posterior marginals for linear predictor and fitted values comput
ed

dic0 = rs0$dic$dic
c(dic0,dic1,dic0-dic1)
## [1] -97.632 -102.880      5.248

# a difference of 5 indicates
# that the smooth fit is somewhat better than the linear fit, but no
t very noteworthy so.

```

6.2 Analysis on nine weighed medians (models G-J)

Fig1b_Nine_medians.r

5.1.2.e

Thu Aug 14 16:27:29 2014

```
rm(list=ls(all=TRUE))
setwd("D:/Data/Aa/_projects/Ctgb")
nature13531.f1 <- read.csv("D:/Data/Aa/_projects/Ctgb/nature13531-f
1.csv")
Data = nature13531.f1
#View(Data)
varnames.ori = names(nature13531.f1)
Data1 = Data
names(Data1)[2:5]= c("y", "x", "species", "Conc.class")
Data = data.frame(Data,Data1)
names(Data)

## [1] "X"                "y_point"            "x_point"
## [4] "Spec."             "Concentration_class" "n"
## [7] "X.1"              "y"                  "x"
## [10] "species"          "Conc.class"        "n.1"

# n = number of plot combinations
# Each point represents the average intrinsic rate of increase of a
species
# over all plots in the same concentration class, whereas
# the size of the point is scaled proportionally to the number of sp
ecies plot combinations on which the calculated mean is based. Binni
ng into classes was performed to reduce scatter noise and aid in vis
ual interpretation. Actual analysis, and the depicted regression lin
e, was performed on raw data (n = 1,459). The regression line is giv
en by  $0.1110 - 0.0374$  (s.e.m. =  $0.0066$ )  $\times$   $\log[\text{imidacloprid}]$  ( $P < 0.0001$ ).
# Dashed lines delineate the 95% confidence interval.
sum(Data$n) # 1459

## [1] 1459

attach(Data)
my.ylab = "Mean species trend"
my.xlab = "ln(Imid. Conc)"
ylim1 = c(-1,0.4) # as in paper
ylim2 = c(-0.25,0.25)
ylim3 = c(-0.6,0.4)
check.outlier = FALSE

specn = species
levels(specn) = 1:nlevels(species)

levels.x = unique(x_point)
```



```

levels.conc = exp(unique(x))
conc = exp(x)
bins = levels(Conc.class)

# summarize per conc class

xm = tapply(x, Conc.class, mean)
nm= tapply(n, Conc.class, sum)
sn = tapply(n,Conc.class, length)
ym.var= tapply(y, Conc.class,var)

# weighted medians per conc class
library(limma)
ym = xm; names(ym)= bins
for (c in bins) ym[c]= weighted.median(y[Conc.class %in% c],w = n[Co
nc.class %in% c ])

Data = data.frame(y= ym,x = xm, n= nm, s = sn, w = 1/ym.var)
attach(Data)

## The following objects are masked from Data (pos = 4):
##
##      n, x, y

print(Data)

##           y      x      n      s      w
## (1.5,2]  0.097983  1.75    7    7   46.90
## (2,2.5]  0.017073  2.25  374  15  239.04
## (2.5,3] -0.014066  2.75  443  15  346.47
## (3,3.5] -0.036495  3.25  170  15  150.81
## (3.5,4]  0.002701  3.75  271  15   87.20
## (4,4.5] -0.056168  4.25   57  11   80.34
## (4.5,5] -0.094018  4.75   72  15   33.54
## (5,5.5] -0.072660  5.25   41  13   40.48
## (5.5,7]  0.004138  6.25   24  13   10.74

# model G
# unweighted
lmod0 = lm(y~1,data=Data)
lmodu=lmod1 = lm(y~x, data = Data)
summary(lmod1)

##
## Call:
## lm(formula = y ~ x, data = Data)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -0.0552 -0.0290 -0.0218  0.0182  0.0780
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)

```

```
## (Intercept)  0.0719    0.0478    1.50    0.176
## x           -0.0233    0.0118   -1.98    0.089 .
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 0.049 on 7 degrees of freedom
## Multiple R-squared:  0.358, Adjusted R-squared:  0.266
## F-statistic:  3.9 on 1 and 7 DF, p-value: 0.0888
```

```
anova(lmod0,lmod1)
```

```
## Analysis of Variance Table
##
## Model 1: y ~ 1
## Model 2: y ~ x
##   Res.Df    RSS Df Sum of Sq   F Pr(>F)
## 1      8 0.0262
## 2      7 0.0168  1   0.00936 3.9 0.089 .
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
#model H
```

```
# weighted by the number of species
```

```
lmod0 = lm(y~1,data=Data, weight= s)
lmods =lmod1 = lm(y~x, data = Data, weight = s)
summary(lmod1)
```

```
##
## Call:
## lm(formula = y ~ x, data = Data, weights = s)
##
## Weighted Residuals:
##   Min      1Q  Median      3Q      Max
## -0.2107 -0.0866 -0.0493  0.0897  0.2613
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  0.0513    0.0472    1.09    0.31
## x           -0.0191    0.0116   -1.66    0.14
##
## Residual standard error: 0.166 on 7 degrees of freedom
## Multiple R-squared:  0.282, Adjusted R-squared:  0.179
## F-statistic: 2.74 on 1 and 7 DF, p-value: 0.142
```

```
anova(lmod0,lmod1)
```

```
## Analysis of Variance Table
##
## Model 1: y ~ 1
## Model 2: y ~ x
##   Res.Df    RSS Df Sum of Sq   F Pr(>F)
## 1      8 0.270
## 2      7 0.194  1   0.0759 2.74 0.14
```

```
# model I
```

```
# weighted by the number of plots
```

```

lmod0 = lm(y~1,data=Data, weight= n)
lmodn = lmod1 = lm(y~x, data = Data, weight = n)
summary(lmod1)

##
## Call:
## lm(formula = y ~ x, data = Data, weights = n)
##
## Weighted Residuals:
##      Min       1Q   Median       3Q      Max
## -0.4034 -0.2232 -0.0988  0.2115  0.4555
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  0.05568    0.03223   1.73   0.128
## x           -0.02151    0.00986  -2.18   0.065 .
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 0.335 on 7 degrees of freedom
## Multiple R-squared:  0.405, Adjusted R-squared:  0.32
## F-statistic: 4.76 on 1 and 7 DF, p-value: 0.0655

anova(lmod0,lmod1)

## Analysis of Variance Table
##
## Model 1: y ~ 1
## Model 2: y ~ x
##   Res.Df  RSS Df Sum of Sq   F Pr(>F)
## 1      8 1.316
## 2      7 0.783 1    0.533 4.76 0.065 .
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

# model J
# weighted by the inverse of variance
lmod0 = lm(y~1,data=Data, weight= w)
lmodw = lmod1 = lm(y~x, data = Data, weight = w)
summary(lmod1)

##
## Call:
## lm(formula = y ~ x, data = Data, weights = w)
##
## Weighted Residuals:
##      Min       1Q   Median       3Q      Max
## -0.2340 -0.1749  0.0421  0.3324  0.4729
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  0.0830    0.0329   2.52   0.04 *
## x           -0.0309    0.0103  -2.99   0.02 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

```

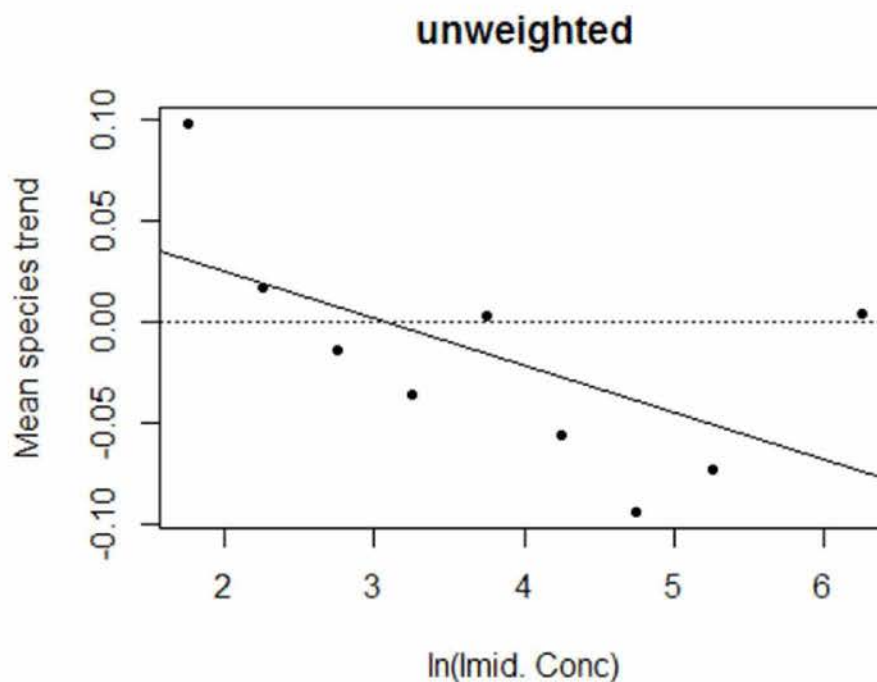


```
##
## Residual standard error: 0.298 on 7 degrees of freedom
## Multiple R-squared: 0.561, Adjusted R-squared: 0.499
## F-statistic: 8.96 on 1 and 7 DF, p-value: 0.0201

anova(lmod0,lmod1)


## Analysis of Variance Table
##
## Model 1: y ~ 1
## Model 2: y ~ x
##   Res.Df  RSS Df Sum of Sq    F Pr(>F)
## 1      8 1.41
## 2      7 0.62  1    0.794 8.96  0.02 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

#par(mfrow=c(4,1))
plot(x,y,pch=20, ylab = my.ylab, xlab= my.xlab, main="unweighted")
abline(0,0, lty = "dotted")
abline(coef(lmodu[1]),coef(lmodu[2]))
```




```
plot(x,y,pch=".", ylab = my.ylab, xlab= my.xlab, main="weight = number of plots")
symbols(x,y, circles = n/2000,inches= FALSE, add= TRUE)
abline(0,0, lty = "dotted")
abline(coef(lmodn[1]),coef(lmodn[2]))
```

weight = number of plots

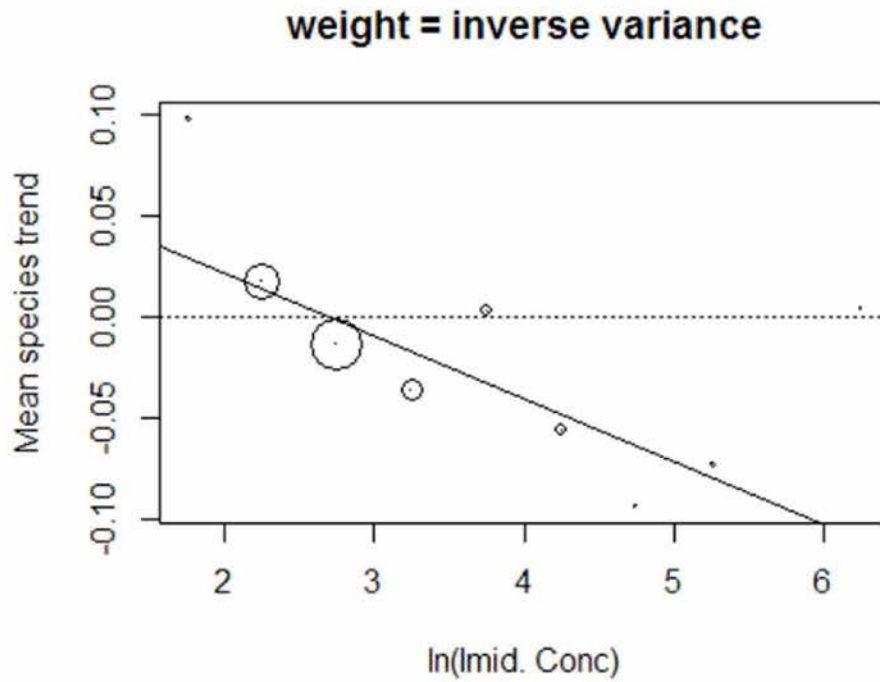


```
plot(x,y,pch=".", ylab = my.ylab, xlab= my.xlab,, main="weight = num  
ber of species")  
symbols(x,y, circles = s/200,inches= FALSE, add= TRUE)  
abline(0,0, lty = "dotted")  
abline(coef(lmods[1]),coef(lmods[2]))
```

weight = number of species



```
plot(x,y,pch=".", ylab = my.ylab, xlab= my.xlab, main = "weight = in  
verse variance")  
symbols(x,y, circles = w/2000,inches= FALSE, add= TRUE)  
abline(0,0, lty = "dotted")  
abline(coef(lmodw[1]),coef(lmodw[2]))
```



End of report