Garden Moth Scheme

Methods for trend analyses

Steve Roberts

GMS recorder and Centre for Biostatistics, University of Manchester

Steve.Roberts@Manchester.ac.uk

SteveRoberts@live.co.uk

V1.0: 16-Sep-25

---

1 Contents

2 Background

This document aims to provide a summary of the data management and statistical methods used in the trend analysis of the 2008-23 GMS data for the UK and Ireland. See other documents for a description of the GMS data and a review of what is available.

We aimed to utilise as much of the data as possible, but the Channel Islands data is too fragmentary and too ecologically different, and we have to leave it out of these analyses.

The 2008-23 period extends from the date where national coverage was largely achieved up to the latest data available. We note that Ireland only joined in 2009.

The first 2 rounds of analyses focus on the main summer scheme (March-Oct). Later analyses will look at the winter species.

3 General Data Cleaning

Data assembly and cleaning script:BuildGMSpart1.R. Data errors listed in GMS data errors.docx

- The field labelled Recorder.Code is actually the Site.Code which links to the habitat database - renamed to Site.Code.
- Exclude records with “Trap not run” in comments. (From 2014 weeks with no data have a dummy empty trap record for an arbitrary date with the fact that the trap was not run added as a comment, before 2014 there are simply no records for these nights)
- We define a *Trap* as the combination of Site.Code and Method - recorders can change traps but should be the same trap for each season.
- The GMS week is not in the database so has to be computed from the date (see code).
- There are a small number of sites with inconsistent grid references (typos?) and in a couple of cases inconsistent VC assignments. We assume the latest data is the correct one. In one case the lates grid reference was clearly incorrect and the previous one is used.
- One record with no date was deleted and a few obvious typos in species codes corrected to be consistent with the recorded scientific and vernacular names.
- Aggregated species: We do not consider it wise to combine these and merge with aggregate records so they are retained as separate species. It seems implausible that all recorders could dissect all these moths for a definitive ID.
- YH-11 used the wrong form in 2008 - removed
- Peppered moth subspecies are allocated the same code. We retain these as a single species as there is significant doubt as to whether the subspecies have been properly recorded.
- There are a handful of trap-nights with both “nothing” and real moths recorded - removed the “nothing” record.
- There are a few records in the database which record moths not on the relevant regional lists (only checked main summer scheme to date). (See data errors document and code in Build2023/BuildGMSpart2.R for details)

    - Moths that were not on the list have been removed (9 records)
    - Moths recorded as single species when the form had aggregated species have been recoded as aggregates (41 records)
    - A very few Copper Underwing were recorded as aggregate when the form only had the Copper Underwing single species – recoded as the single species (3 Irish records only)

4 Inclusion/exclusion Criteria

The aim here was to be as inclusive as possible to value the data collected and the recorders' efforts.

We define 3 datasets here:

4.1 Full GMS data

Code in BuildGMSpart1.R

Everything after cleaning (All.Rd)

4.2 Main National Summer scheme (excluding CI)

Code in BuildGMSpart1.R

- Summer scheme - all records in the 36 weeks of recording
    - Defined as first Friday-3 days to last Friday (36 weeks later) + 3 days
- National coverage. 2007 was nominally the start of the national rollout but 2007 data rather patchy. 2008 is much better but Ireland and CI only came on board in 2009. So we currently define the national scheme from 2008 to the latest available (2023). There is an argument for starting from 2009.
- Exclude Channel Islands. As a matter of principle, we would like to include as many of the recorders as possible. However:
    - It would be unusual to include the CI in a UK and Ireland study as they are geographically rather distant and have a different climate and fauna.
    - There are very few traps with more than a couple of years data. For any analysis they would have to be pooled with another region, which would be implausible.
    - Therefore, regrettably, we have to exclude the CI from the main analyses. Maybe at some stage we could consider a CI-specific analysis.

4.3 Main Season Trend analysis subset

Code in BuildGMSpart2.R

- For the trend analyses (second stage – see below) we Include all traps with 3 or more seasons data after start of national scheme (2008). There needs to be at least 2 years of data to be informative for trends, but very little information from short sequences and fitting many barely informative traps causes numerical issues in the analysis. We have currently chosen 3 years as a compromise to include as much data as feasible.
- Exclude trap/years with <25 weeks of records (very few)

4.4 Winter scheme (excluding CI)

Code in BuildWinter.R

- Winter scheme covers 16 (or occasionally 17) weeks Nov-Feb starting the week after the end of the summer scheme and finishing the week before the start of the next summer scheme. We consider this to be weeks 37-52 (or 53) of the previous main season. We create a dataset containing just the winter data, but for analysis we create an expanded dataset which includes the main season for the traps included in that year's winter season. We note that many recorders switch traps between the summer and winter schemes, and we only include data from same trap (not the same garden). This expanded dataset allows us to estimate abundances and phenology for species that fly across both the main and winter scheme periods.

4.5 Winter Season Trend analysis subset

Code in BuildWinter.R

- For the trend analysis we again limit to traps with 3 or more seasons of data.
- We exclude traps with <8 nights of recording.

5 Data Assembly

5.1 Main Summer Series dataset

This is the cleaned dataset created in *buildGMSpart1.R* and the cleaning and inclusion criteria are as above.

- Saved in Main.Rd with Trap.main, TrapWeeks.main and TrapYear.main

5.2 Main Winter Series dataset

This is the cleaned dataset created in *buildWinter.R* and the cleaning and inclusion criteria are as above.

- Saved in winter, Trap.winter, TrapYear.winter,TrapWeeks.winter

5.3 Augmented Winter Series dataset

This is the augmented dataset created in *buildWinter.R* which includes data from the main scheme where the same trap was used.

- Saved in winterPlus, Trap.winterPlus, TrapYear.winterPlus,TrapWeeks.winterPlus,

5.4 Trend Analysis dataset - main season

This is the cleaned dataset created in *buildGMSpart2.R* and the cleaning and inclusion criteria are as above.

- Saved in Trend.Rd with datasets: Trend, Trap.trend, TrapWeeks.trend and TrapYear.trend

5.5 Trend Analysis dataset - winter season

This is the cleaned **augmented** dataset created in *buildWinter.R* and the cleaning and inclusion criteria are as above.

- Saved in WinterTrend.Rd with datasets: Trend, Trap.trend, TrapYear.trend and TrapWeeks.trend

5.6 Species List

A dataset/file detailing which moths were on the list for each region/year

Code in ReadRecordingForms.R

- Read all the recording forms for 2007-2024
- A few missing forms can be assumed to be the same as the adjacent years (ie no changes over the years before and after).
- This leaves CI 2009 and most of the 2007 forms. We use the 2008-24 data to impute a CI 2009 form based on what was recorded – an under-estimate. 2007 we ignore for now.
- For each year/Region create a dataset as a list for each region+core containing a data frame with rows for each moth and columns with a T/F entry for whether that moth is on the regional (or core) list.
- In a final data frame we collate a similar data frame indicating those moths on all the regional lists for each year – an effective core list.
- The winter lists were common across all regions and is saved in an analogous winter dataset

- These lists are written to an Excel file SpeciesFile.xlsx in the working directory and saved as SpeciesList.Rd in R format.

5.7 Trap datasets

Code in buildGMSpart2.R, BuildWinter.R

Auxiliary files to the main and trend datasets containing characteristics of the included traps:

- Method: Trap description
- Year.First: First year recorded
- Years: Years recording
- MeanMoths, MeanSpecies: Mean moths and species per year

5.8 Trap-Year datasets

Code in buildGMSpart2.R, BuildWinter.R

Auxiliary files to the main and trend datasets containing summary data from the included traps with one row per trap-year:

- Nights: Nights recorded
- Moths: Total Moths recorded
- Species: Total Species recorded
- Gridref: Grid reference
- Lat, Long: Location latitude and longitude

5.9 Trap-Week datasets

Code in buildGMSpart2.R, BuildWinter.R

Auxiliary file to the main and trend datasets containing one row per trap-week where trap was run. Simply used to identify when each trap was run.

6 Species Included in Trend analyses

- We select all the 180 species recorded in the main summer scheme for all regions from 2010 onwards where at least 2000 moths were recorded. Most of these species were recorded in nearly all regions in 2008 and 2009. We exclude 12 species where the flight season extended beyond the main summer season and 2 species that did not permit reliable trend estimates, leaving 166 species analysable from the main summer scheme datasets.
- Additionally we analysed 12 species from the (augmented) winter scheme where there were at least 2000 moths recorded, and where the species had been on the recording list since 2010. This number excluded a few very common moths with a flight season covered by the main scheme, but had winter records and one species which failed to yield reliable trend estimates.

7 Species-level Data Preparation for Analysis

[function `MothSetUp` in `SetupTrendData.R`]

- For each species we create datasets `moth.main` and `moth.trend` with just the records for that species in the `Main` and `Trend` datasets

- Add zero records where trap was run but moth not recorded (taken from data in `Trapweeks.main` and `TrapWeeks.trend`)
- Remove any entries where the moth was not on the regional list for the year (very few) (data in `SpeciesLists`)
- Remove all entries in `moth.trend` dataset for traps where the moth was never recorded in any year (uninformative and numerically problematic)
- Creates datasets Mothxxx.Rd for main scheme species and MothxxxW.Rd for winter scheme species (saved in Working directory), where xxx in the ABH code

8 Primary Analysis

The primary analysis uses a 2-stage approach, conceptually similar to that used for UKBMS and RIS data, but with technical differences due to the different data structure.

- Stage 1 estimates an annual abundance for the species in each trap and each year
- Stage 2 estimates the trend in these abundances across the 16 years

The primary unit of analysis is a combination of the garden and the trap type, referred to as the “trap”. That is, we assume that the capture rate for a garden may change if the type of trap or bulb changes. (The trend analysis adjusts for trap and essentially estimates trends within traps and therefore does not need to further adjust for garden/trap characteristics).

8.1 Stage 1: Estimating the annual abundances

(function `ComputeAUC` in `AbundanceModels.R`)

1. For this stage we used all the main summer (non-CI) data – the main dataset.
    - This ensured we made maximum use of the data provided by the scheme participants, including those who only participated for 1 or 2 years
2. For each year we fit the number of moths observed in model including a B-spline function based on the *calendar day* of year the trap was run to model the phenology and a trap-specific parameter for each trap. These are fitted using a generalised linear model with a negative-binomial distribution.
    - For most species we used a B-spline with 10 degrees of freedom which was deemed sufficient to capture the shapes of the flight period curves. In a few cases this proved excessive and led to artefacts and 8df (or in one case 6df) splines were fitted.
    - We noted that there are small numbers of records for many species that were outside the usual flight period (see individual species reports). In general we retain these records and they have negligible impact on the abundance estimates. In a few cases the phenology plots are rather distorted by these outliers and in these cases (currently just 3) we exclude these observations (see `Exclude.xlsx` for a list)
    - We also fitted a Poisson model, but the data are highly over-dispersed so we prefer the negative-binomial model. In most cases the abundances are similar, but there are cases where the two models differ non-trivially.
    - It is worth noting that a simple linear interpolation of missing weeks and computation of the abundance in each trap/year using a Trapezoidal rule gave very similar results to the Poisson spline model – the GMS data has very few missing weeks.
    - We note this assumes each site has the same phenology in each year (but this may vary between years). This is the same assumption used in the published analyses of the RIS data

    - We use the R package `glmmTMB` as this proved the faster and more reliable of the `R` glm packages for this data.
    - Preliminary analyses using a different model showed that once we allow for overdispersion with a negative-binomial model the fitted number of zeros was close to that observed so we didn't consider including zero-inflation in the modelling
3. For each trap (in each year) we estimate the fitted model for each day and sum to get an abundance for that year.
4. For each year we estimate the phenology curve (at day resolution) normalised to have unit area.
    - In these analyses we simply display these phenology curves, future analyses could use these annual curves to look for trends in flight periods over time in a second stage analysis similar to that used for abundance.

8.2 Stage 2: Estimating Trend

1. (Function `FitAUC` in `AbundanceModels.R`) Based on the trap/year abundances we fit linear models with separate intercepts for each trap to the data in the trend dataset (excluding any traps that did not ever record the species) assuming a negative-binomial distribution:
    a. With a categorical variable for year to get estimates for the overall relative abundance for each year (fitted estimates with their Wald standard errors (computed on the link scale and back-transformed) for an arbitrary trap divided by the estimate for 2008 – presented as estimates with 95%CI.
    b. With a linear variable for year to estimate the trend slope (change per year) with 95% Wald-based CI. This is subtracted from 1 for presentation and expressed as a %change per year.
- We also fitted models assuming a Poisson distribution – generally the estimates and CI were similar, but the data did show overdispersion, and in some cases the estimates did differ non-trivially so we present the negative-binomial results.
- Additionally, we tried allowing for autocorrelation with an AR1 structure. In nearly all cases this gave identical slope estimates – the overdispersion seeming to be sufficient to absorb any autocorrelation effects.

2. (Function `AddBootstrap` in `AbundanceModels.R`). The second stage was then bootstrapped (resampling Traps, 1000 samples) to obtain revised 95%CI for the slope estimates.
    - This follows the precedent from the published RIS analyses, although they also resampled the first stage.
    - It is likely that this resampling is conservative (giving over-wide confidence intervals) as some of the structure of the dataset is lost with fewer joins across traps in this very fragmented dataset – the resamples would be expected to have less statistical power than the full dataset.
    - The alternative of resampling trap-years was considered but this is likely to be over-conservative as the, already short, within-trap series will cover shorter periods in the resamples (they cannot be longer).
    - We note that for RIS-type data with fewer longer series these issues would be expected to be trivial, but here, and maybe for UKBMS transect datasets, it is not known if bootstrapping gives reliable Cis. This could be explored in a future simulation study.
    - However as the bias is towards over-wide CI, and as we generally have good statistical power, we choose to present these bootstrapped estimates and retain broad compatibility with RIS methodology.
    - We did also bootstrap a Poisson fit to AUCs derived from a Poisson model. This showed:

  - For most moths the CI were far too narrow from this model and some form of bootstrapping would be essential to get reasonable fits
  - The Poisson-Poisson approach was seriously biased for a few moths

9 North-South Trend differences

Based on the observation in the 2021 State of UK moths report [Fox R, Dennis EB, Harrower CA, Blumgart D, Bell JR, Cook P, Davis AM, Evans-Hill LJ, Haynes F, Hill D, Isaac NJB, Parsons MS, Pocock MJO,Prescott T, Randle Z, Shortall CR,Tordoff GM,Tuson D & Bourn NAD (2021) The State of Britain's Larger Moths 2021. Butterfly Conservation,Rothamsted Research and UK Centre for Ecology & Hydrology, Wareham, Dorset, UK.] we looked at differences in trends between the Northern and Southern halves of the UK. For consistency we used the same criteria as the previous report and defined the two halves as above and below the 450km N line of the Ordnance Survey National Grid (53.7920 latitude).

The first stage of the analysis was performed separately for the two halves of the UK, thus allowing the phonologies to differ between the regions and computing the abundance for each trap-year. The data were then combined for the second, trend-fitting, stage, and a year x region interaction term was added to the negative-binomial model, parameterising the model to obtain trends estimates for the North and South. This second stage was bootstrapped (resampling traps, 1000 resamples) to obtain improved confidence intervals and P-value (the latter based on the number of samples with positive or negative differences in the slopes).

As this analysis can only be performed for species that occur in sufficient numbers in both halves of the country we only included species that had at least 1000 moths recorded in both halves.

10 Presentation: Individual species reports

[`OneMothPageV3.qmd` (older extended version in `OneMothPlus.qmd` with more diagnostics) Quarto/Typst markdown code used to mass-produce the reports]

The one-page summary includes:

1. Some notes giving a little context based on the Moth atlas (RIS data) and a summary of the trends observed [Notes taken from `SpeciesNotes.xlsx`]
2. A table for the Main and Trend datasets showing the number of records, total number of moths recorded and the number of Traps and Trap-years the species was recorded. The numbers of traps and trap-years are expressed as a percentage of the numbers where the species was on the relevant recording list. Note that for the trend dataset we select only traps which have recorded the species. (Note the trend dataset is restricted to traps that recorded the species, so always has 100% of traps recording the species)
3. Distribution maps showing (red closed symbols) traps recording the species and (open blue symbols) traps not recording the species. These plots are presented for 3 time periods, and to allow for differing numbers of years each trap may have been in the scheme *we currently define a species presence as being recorded on an average of at least once per year in the trap in that period*. Distributions only compared visually over time as the trap locations change and so trends in distributions cannot be reliably quantified.
4. A plot of the phenology for each year derived in stage 1. These are either
    a. coloured blue through red for successive years if there were no North-South differences or

    b. Coloured red and blue for North and South in each year.
5. A plot of the relative abundances over the time period, normalised to the 2008 abundance.
    a. The annual abundances derived in stage 1 with the fitted 95%CI (not bootstrapped). Where there were significant North/South differences both are plotted, otherwise a pooled national estimate is shown.
    b. The fitted trend line(s) from stage 2
    c. Where there are no significant North/South differences, text giving the slope of the trend line, expressed as % per year with the bootstrap-derived 95%CI and a note where this reached nominal statistical significance (with the usual caveats about arbitrary P<0.05 criteria)
    d. Where there are North/South differences text shows the slopes for the North and South with 95%CI and a P-value for the North/South difference in trend.

We characterise the number of moths seen and their distributions with the following terminology:

Number of moths recorded		**Number of gardens recording**	
>100k	Extremely Common	>95%	Ubiquitous
50-100k	Very common	>67%	Widespread
10—50k	Common	33-67%	Regional (noting where)
2-10k	Uncommon		Local (if no obvious range)
<2k	Excluded	<33%	Regional
			Very local

Although we will diverge from a strict characterisation where this seems appropriate and for moths using the winter data we use values ¼ the size (i.e. roughly half the duration and half the traps).

11 Associations of species Traits with abundance trends

As an initial exploration the summary document plots the rate of change of the GMS abundances against various species traits from the Cook *et al* Traits database [add ref]. These are exploratory plots and need some more work!

We also use a file, `Families.xlsx,` derived from the UK checklist to assign species to families.

As we found nothing of note, these are not reported but are available in the technical report of main season trends.

12 Alternative Analyses Explored

An alternative approach fitted a single stage model to the full species data, with a single spline model for phenology (using day of year) for all traps and years. This gave very similar trend estimates for nearly all the moths. As well as the statistical advantage of a single-step approach, the hope was that this model would have the flexibility to explore things like the regional variation in phenology. However, these models proved tricky to fit reliably and extension beyond the simple common phenology was not feasible. Therefore, these models were used as a sensitivity/validation analyses of the 2-stage approach described above.