Stabilising sparse-data trends in the National Plant Monitoring Scheme: a scalable Bayesian model for plant abundance*

Oliver L. Pescott† Robin J. Boyd
UK Centre for Ecology & Hydrology, Wallingford, UK, OX10 8BB

2025-08-19

Abstract

The UK National Plant Monitoring Scheme produces around 600 annual species within-habitat trends that combine information on occupancy and positive cover. These trends underlie developing government indicators. We give a full specification of the Bayesian model and priors, and document changes introduced in the 2025 run (covering 2015–2024) relative to the 2024 run (2015–2023) to stabilise trends for sparsely recorded species. The principal change replaces first-order random-walk temporal priors with stationary autoregressive (AR(1)) priors; we also tighten several detection and cover priors. Future work will address non-response bias in NPMS survey returns.

1 Introduction

The UK National Plant Monitoring Scheme (NPMS) was launched in 2015 following several years of design and field trials (Pescott et al., 2019b). Following this, Pescott et al. (2019a) developed a model for submitted data in order to create within-habitat species-level trends. This model coupled an occupancy-detection submodel for plot occupancy with an interval-censored Beta distribution submodel for non-zero cover, allowing for the tracking of the means of potentially zero-inflated cover estimates over time. Subsequently this model was developed in various ways for the UK Biodiversity Indicator (UKBI) suite. The main developments were extensions of the detection submodel and refinements to the priors designed to stabilise the model when run across ~600 species-habitat combinations. This report details the current model in full, whilst also noting key developments introduced in 2025 (i.e. for the 2015–24 indicators) that have, in some cases, resulted in changes to the habitat-level multi-species indicators presented under UKBI C7 "Plants of the wider countryside".

2 Model description

The following algebra describes version 5.0 of the model (Appendix 1), used to produce the NPMS species-habitat trends underlying the UKBI C7 "Plants of the wider countryside" indicator, 2015–24. The section "Changes from previous years" below describes the differences from the previous version of the model used for the C7 indicator published between 2020 and 2024 (see here for archived versions of these).

2.1 Indexing

Define the following index sets and mappings:

*Cite as: Pescott, O.L. & Boyd, R.J. 2025. Stabilising sparse-data trends in the National Plant Monitoring Scheme: a scalable Bayesian model for plant abundance. Unpublished report to Defra. UKCEH: Wallingford.

†olipes@ceh.ac.uk

$$\begin{aligned}
&\mathcal{A} = \{1,\dots,V\}, && \text{all visits}\\
&i : \mathcal{A} \to \{1,\dots,N\},\ a \mapsto i(a), && \text{plot index for visit } a\\
&j : \mathcal{A} \to \{1,\dots,Y\},\ a \mapsto j(a), && \text{year index for visit } a\\
&\mathcal{A}_{\text{pos}} = \{a \in \mathcal{A} : \text{visit } a \text{ had positive cover}\}.
\end{aligned}$$

Importantly, although the NPMS recommends that plots are visited twice a year, $V \neq N \cdot Y \cdot 2$. This is because many plots are not sampled every year, or on both potential within-year sampling events.

2.2 Plot occupancy submodel

Recalling that each visit a maps to a specific site-year pair [i.e. $a \mapsto (i(a), j(a))$], true occupancy (presence/absence) is a binary variable $z_{i(a),j(a)} \in \{0, 1\}$, which is assumed to follow a Bernoulli distribution:

$$z_{i(a),j(a)} \sim \text{Bernoulli}(\psi_{i(a),j(a)}),\ \text{logit}(\psi_{i(a),j(a)}) = m_{j(a)},$$

where $\psi_{i(a),j(a)}$ is the probability that the modelled species occupies plot i in year j, and $m_{j(a)}$ is a year-specific intercept shared across plots. Whether the species was detected on visit a, conditional on occupancy, is modelled as:

$$y_a \sim \text{Bernoulli}(z_{i(a),j(a)} \cdot p_a),$$

where $z_{i(a),j(a)}$ is the true occupancy for that visit and $p_a \in (0, 1)$ is the detection probability. The detection probability is modelled as:

$$\text{logit}(p_a) = g_0 + g_1\, c^*_a + g_{2,x_{2[a]}} + g_{3,x_{3[a]}}$$

where g_0 is an intercept, g_1 is the coefficient for plant cover c^*_a, g_2 is the coefficient for the effect of different levels of grazing or mowing of the plot, and g_3 is the coefficient for the level at which a surveyor participates in a scheme (i.e. an index of expertise).

A challenge here is that we are attempting to use plant cover as a covariate influencing plant detection, but if a surveyor overlooks a species (which may often be due to low plot cover), we have no empirical information on cover to use in the detection model. To overcome this limitation, an indicator variable δ_a is used to switch the covariate between observed and predicted cover where necessary (in fact, the empirical cover estimates are on the ordinal Domin scale, so an additional step, absent from the algebra below, is the conversion of predicted cover $\hat{c}_a$ onto that scale prior to integrating it with c_a to form c^*_a; this is given in the JAGS code in Appendix 1). Formally,

$$\delta_a = \begin{cases} 1, & \text{if the empirical cover estimate } c_a \text{ exists,}\\ 0, & \text{otherwise.} \end{cases}$$

Then,

$$\begin{aligned}
c^*_a &= \delta_a\, c_a \ + \ (1-\delta_a)\, \hat{c}_a,\\
\hat{c}_a &\sim \text{Beta}\big(\alpha_{j(a)},\, \beta_{j(a)}\big)\ T(10^{-5},\, 1-10^{-5}).
\end{aligned}$$

Note here that whilst visits map to specific plot-year pairs, unobserved predicted cover values are drawn from a Beta distribution parameterised at the year level only (see section “Non-zero plot cover submodel” below for information on how the Beta distribution is parameterised from observed data). As elsewhere, the Beta distribution truncation is used for numerical stability within JAGS.

2.3 Non-zero plot cover submodel

Recall that $\mathcal{A}_{\text{pos}} \subset \mathcal{A}$ indexes those visits that recorded some (non-zero) cover. Surveyors report cover values according to the Domin scale, an ordinal scale for which proportional cover intervals can be interpreted (e.g. Pescott et al., 2016). That is to say, each positive cover value for visit a in $\mathcal{A}_{\text{pos}}$ has associated lower and upper bounds $[L_a, U_a]$, where $L_a > 0$, $U_a < 1$, and $L_a < U_a$; occasional presences lacking cover data can also be included by mapping to the full interval $10^{-5} \le \tilde{c}_a \le 1 - 10^{-5}$. The latent cover $\tilde{c}_a$ thus has the constraint $L_a \le \tilde{c}_a \le U_a$, and contributes the likelihood factor $\Pr(L_a \le \tilde{c}_a \le U_a)$ towards inference on the Beta distribution shape parameters $(\alpha_{j(a)}, \beta_{j(a)})$ determining the latent distribution of positive covers within years. Note that (α_j, β_j) are ultimately parameterised via the standard transformations $\alpha_j = \mu_j \phi_j$ and $\beta_j = (1 - \mu_j) \phi_j$, where μ_j is the Beta distribution mean and ϕ_j its precision for year j.

The yearly Beta shape parameters (α_j, β_j) are then used to draw (uncensored) latent true covers per plot-year (i, j) combination, i.e. $\tilde{c}_{i,j}$, and these are multiplied by the latent true occupancy state estimates, $z_{i,j}$, to provide a zero-inflated cover random variable $C_{i,j}$; that is, $C_{i,j} = \tilde{c}_{i,j} z_{i,j}$. The yearly mean of $C_{i,j}$ across plots is then used to construct each species-habitat trend in abundance, i.e.

$$\bar{C}_j = \frac{1}{N} \sum_{i=1}^{N} C_{i,j}.$$

The species-habitat mean-cover values $\bar{C}_j$ and their posterior uncertainties (95% credible intervals) are updated annually on the NPMS website (see https://www.npms.org.uk/trends).

2.4 Prior distributions

The model is fitted in a Bayesian framework, and so all parameters and missing or latent data require prior (probability) distributions.

2.4.1 Plot occupancy submodel priors

2.4.1.1 Year-specific occupancy intercept The year-specific occupancy intercepts m_j are modelled as a first-order autoregressive process. Let

$$\begin{aligned} \alpha_m &= \operatorname{logit}(\mathbb{E}[\psi_{i,j}]), \\ \sigma_m &\sim t_+(0, 1, 3)\, T(0, 0.5) \\ \tau_m &= 1/\sigma_m^2, \\ \rho_m &= \operatorname{Uniform}(0, 1). \end{aligned}$$

Then,

$$\begin{aligned} m_1 &\sim N(\alpha_m,\ \tau_m^{-1}), \\ m_j &\sim N(\alpha_m + \rho_m\,(m_{j-1} - \alpha_m),\ \tau_m^{-1}), \quad j = 2, \ldots, Y. \end{aligned}$$

Here α_m is the (logit-scale) prior long-run mean occupancy (common to all plots and years under the AR(1) prior); $m_{j-1} - \alpha_m$ denotes the deviation from the long-run mean occupancy in year $j - 1$; $\rho_m \in (0, 1)$ is the temporal autocorrelation coefficient; and σ_m^2 is the year-to-year variance controlling how much m_j can change from one year to the next (the inverse of this is τ_m, the precision). Note that the AR(1) normal distributions are also truncated on the logit scale for numerical stability (Appendix 1). In addition to those already specified, the hyperprior for long-run mean occupancy on the logit scale is given by $\alpha_m = \operatorname{logit}(\theta)$, where $\theta \sim \operatorname{Beta}(2, 18)\, T(10^{-5}, 1 - 10^{-5})$. This is visualised in Figure 1.

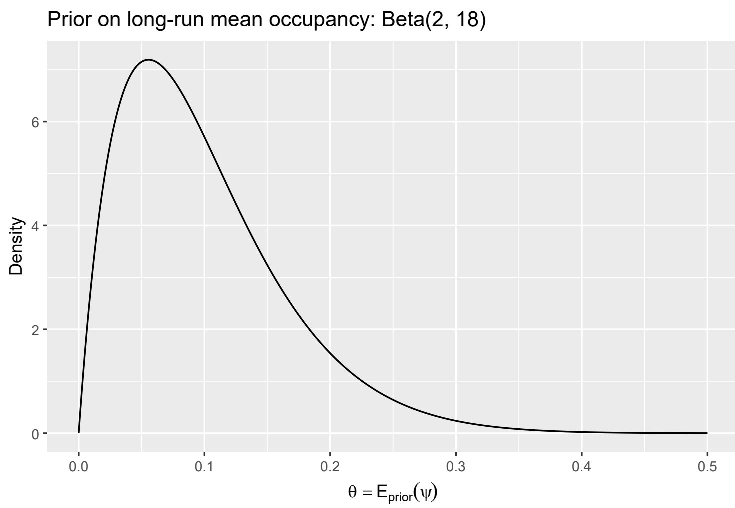

Figure 1: Prior density on grand mean occupancy.

Finally, note that the truncated Student's t distribution specified above as $t_+(0, 1, 3)$ has a mean of zero, precision of one, and three degrees of freedom; see Plummer (2017) for more information. Figure 2 shows the truncated version of this distribution used in this instance.

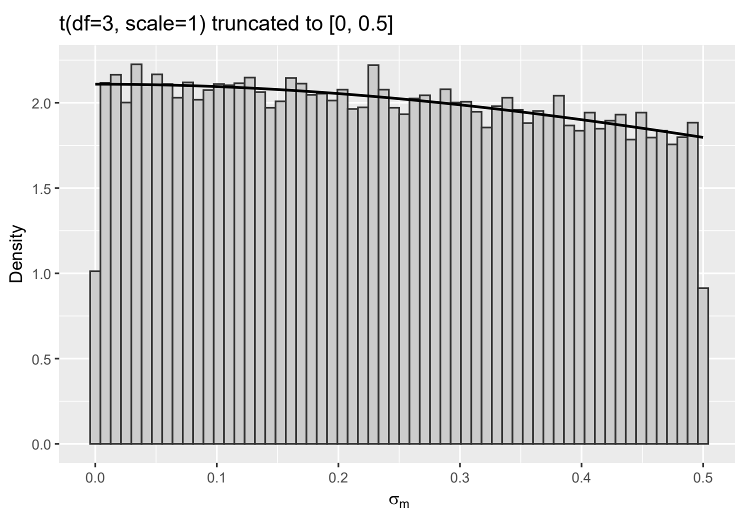

Figure 2: Prior density on innovation standard deviation for year-to-year logit-scale occupancy. The curve is generated from the theoretical probability density function, whereas the histogram is generated using the R package crch.

These hyperpriors create an implied prior for ψ_1 as shown in Figure 3.

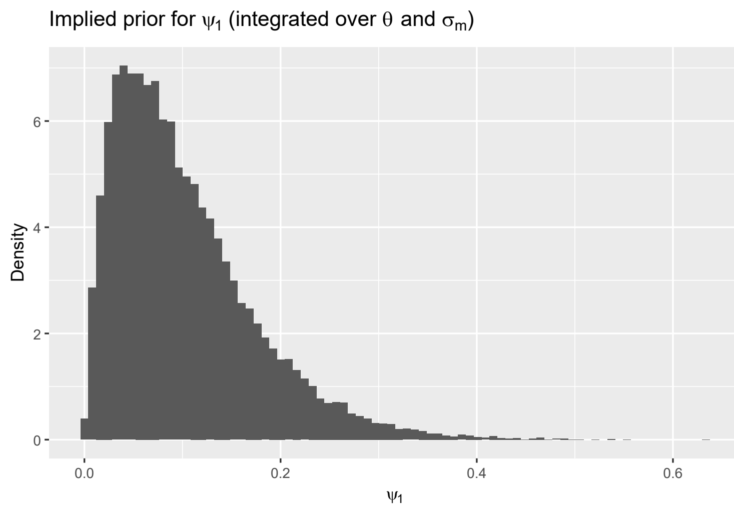

Figure 3: Implied prior density for initial occupancy obtained through combining priors on grand mean occupancy and its logit-scale standard deviation.

2.4.1.2 Detection parameters The following priors are used for the intercept and logistic regression coefficients in the detection submodel. For the intercept: $g_0 = \text{logit}(\pi)$, where $\pi \sim \text{Beta}(5, 5)\ T(10^{-5}, 1 - 10^{-5})$; see Figure 4.

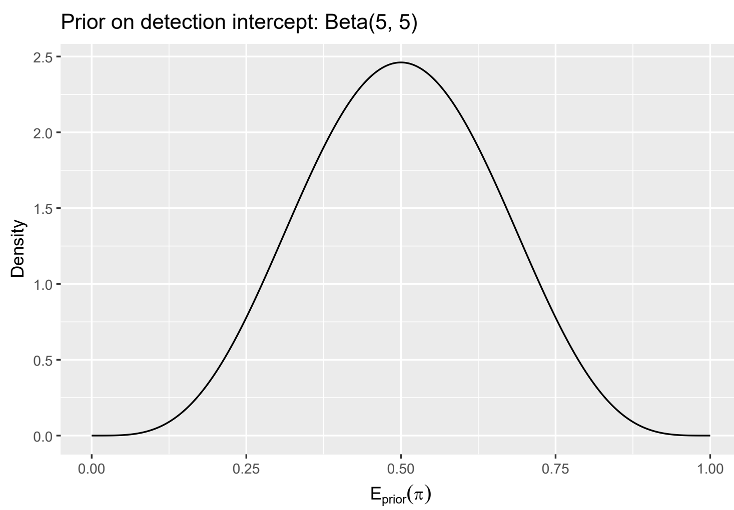

Figure 4: Prior density on the intercept of the detection model.

And for the cover slope: $g_1 \sim \mathcal{N}(1, 1/0.5)\ T(0, \infty)$, note here that within the JAGS code the normal distribution is parameterised using its precision τ, and $\sigma^2 = 1/\tau$. This prior implies that the relationship between cover and detectability is expected to be positive. The priors on the coefficients for grazing/mowing and surveyor level are slightly more involved: they use hierarchical priors with hyperpriors on a shared mean and variance, which provide partial pooling (i.e. shrinkage of level effects towards the group mean). Given that they are identical in structure, we provide the prior structure for g_2 alone below:

$$g_{2,l} \mid \mu_2\, \sigma_2^2 \sim \mathcal{N}(\mu_2,\, \sigma_2^2), \quad l = \{1, \ldots, L_2\}$$
$$\mu_2 = \text{logit}(\pi_2), \quad \pi_2 \sim \text{Beta}(1,\, 1),$$
$$\sigma_2 \sim t_+(0,\, 1,\, 3)\; T(0,\, \infty),$$

where $t_+(0,\, 1,\, 3)\; T(0,\, \infty)$ denotes a half-Student's t prior (df = 3, precision = 1), and l indexes the levels of the categorical variable.

2.4.2 Non-zero plot cover submodel priors

The prior on mean positive cover is also modelled as a first-order autoregressive process. Let $m_{C,j} = \text{logit}(\mu_{C,j})$. This is initialised and evolves via:

$$\mu_{C,1} \sim \text{Beta}(1,\, 1), \quad \alpha_C = \text{logit}(\mu_{C,1}),$$
$$m_{C,1} = \alpha_C, \quad m_{C,j} \mid m_{C,j-1},\, \alpha_C,\, \sigma_C,\, \rho_C \sim \mathcal{N}\big(\alpha_C + \rho_C\{m_{C,j-1} - \alpha_C\},\, \sigma_C^2\big)\,,$$

where Beta(1, 1) is uniform on the cover scale, and with hyperpriors

$$\rho_C \sim \text{Uniform}(0,\, 1), \quad \sigma_C \sim t_+(0,\, 0.5,\, 3)\; T(0,\, 0.5).$$

Finally, $\mu_{C,j} = \text{ilogit}(m_{C,j})$. The prior on the year-specific precision of the Beta distribution, ϕ_j, is an independent Pareto prior:

$$\phi_j \sim \text{Pareto}(k,\, c), \quad f(\phi_j) = k\, c^k\, \phi_j^{-(k+1)} \text{ for } \phi_j \geq c,$$

with $k = 5$ (shape) and $c = 0.5$ (scale) producing heavy tails, i.e.

$$\phi_j \sim \text{Pareto}(5,\, 0.5), \quad j = 1, \ldots, Y.$$

3 Changes from previous years

3.1 The motivating problem

Between 2020 and 2024 the National Plant Monitoring Scheme model in use (v. 4.5) used first-order random-walk priors for the year-specific occupancy intercepts and on yearly mean positive cover: that is, $m_j \sim \mathcal{N}(m_{j-1},\, \sigma_m^2)$. Here, in the absence of data, each year's parameter is centred on the previous year's value, leading to a non-stationary time-series whose uncertainty accumulates over time. Inspection of the initial results from this model in 2025 revealed that some data-sparse species-habitat combinations exhibited mean-cover values ($\bar{C}_j$) that were notably different to earlier outputs, and which, moreover, were deemed ecologically unlikely. Figure 5 below demonstrates the issue.

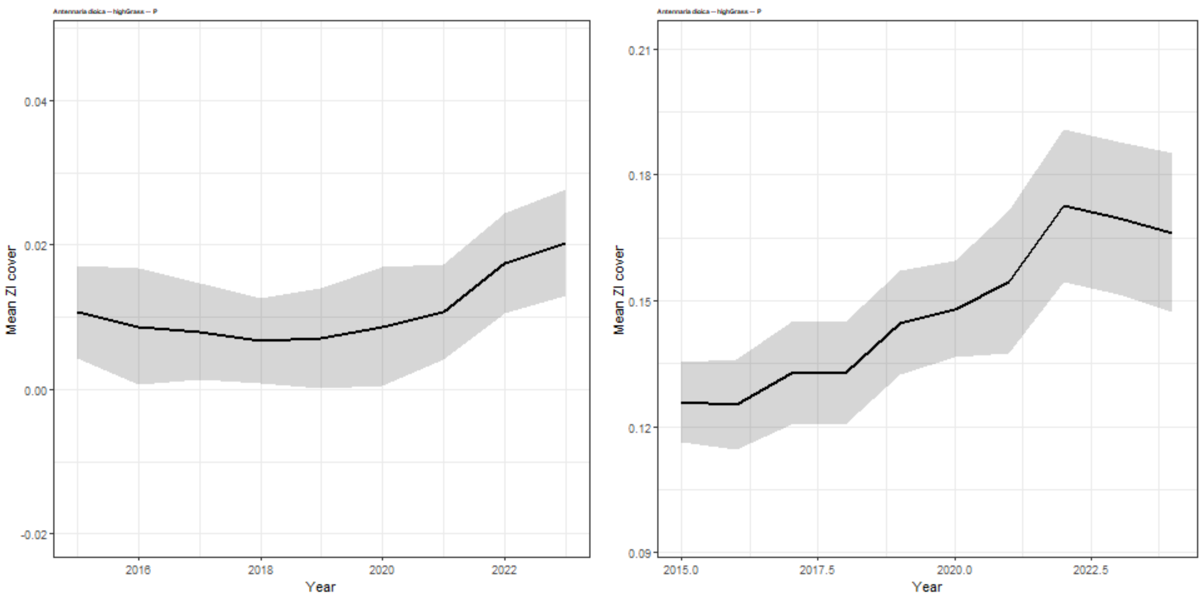

Figure 5: Trends for *Antennaria dioica*, Upland grassland, using model version 4.5. 2015–2023 (left) and 2015–2024 (right). Note the y-axis shift.

Here, the previous year's time trend (i.e. for 2015–2023) for *Antennaria dioica* in upland grassland gave a typical mean (zero-inflated) cover of 1% and the jump to a mean of ~15% for the period 2015–2024 was clearly problematic (Fig. 5), especially given that inspection of the raw data showed that no additional observations of *Antennaria*, or amendments to historic data, had been made in 2024. The use of the updated model (v. 5.0) returned the trend to sensible values, in line with the earlier time-series (Fig. 6). Similar cases were found for other data-poor species.

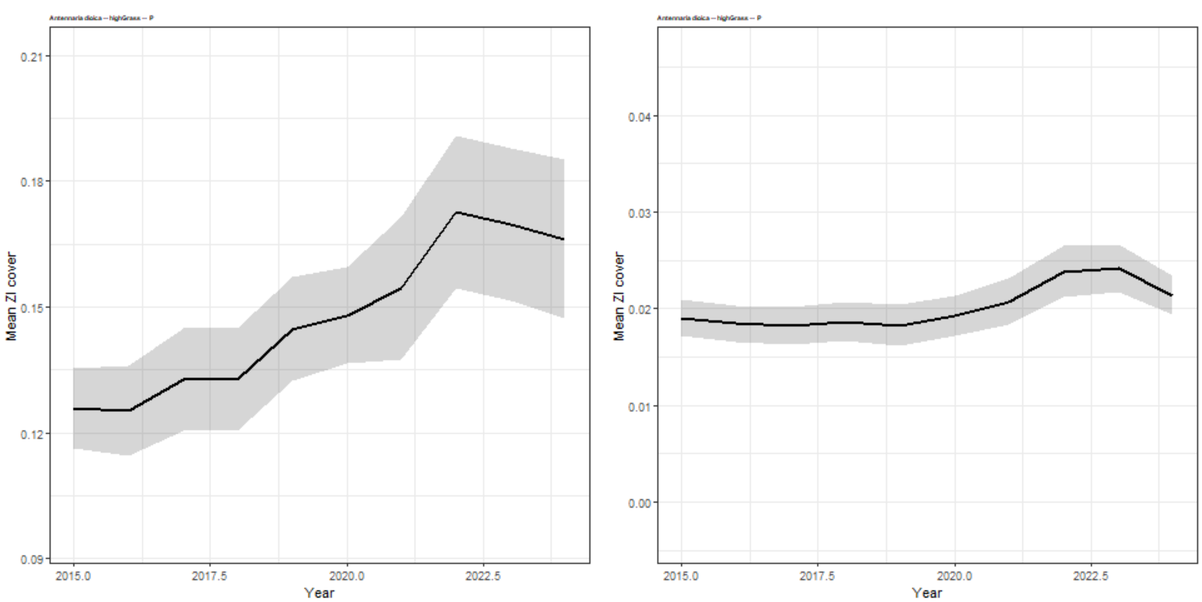

Figure 6: Comparison of model versions for *Antennaria dioica*, Upland grassland, 2015–2024. Model version 4.5 (left) and 5.0 (right). Note the different y-axis scales.

3.2 Justifying the changes

Whilst the adoption of a new model form appeared to produce more sensible results across both data-rich and -poor species (recall that all current trends can be viewed at https://www.npms.org.uk/trends), the changes also require theoretical justification. The following sections provide this.

3.2.1 Random-walk versus autoregressive priors

In model version 4.5 we placed first-order random-walk (RW1) priors on the year-specific intercepts $m_j = \text{logit}(\psi_j)$ and $m_{C,j} = \text{logit}(\mu_j)$: $m_j \mid m_{j-1} \sim \mathcal{N}(m_{j-1}, \sigma^2)$. This non-stationary prior accumulates variance with time and the joint prior over the full time-series depends on the number of years. When the terminal year Y contains no detections and no positive-cover observations, the likelihood contributes little information for m_Y. Under RW1, extending the series by one year can therefore re-centre the entire latent trajectory to satisfy the smoothness constraint: ultimately there is a penalty on successive annual differences implied by the RW1 prior, keeping adjacent years close. This can yield an apparent intercept translation (i.e. shifts up or down the y-axis) in ψ_j and μ_j (and hence in their product $\bar{C}_j$) across all years. The replacement of the RW1 with stationary AR(1) priors $m_j \sim \mathcal{N}(\alpha + \rho(m_{j-1} - \alpha), \sigma^2)$, which shrink towards a long-run mean α_m, limits variance growth and minimises the influence of a data-poor terminal year, preventing wholesale shifts in the trajectory while still allowing year-to-year deviations when data are informative. The following plots demonstrate some aspects of this behaviour using a simulated example (Figure 7).

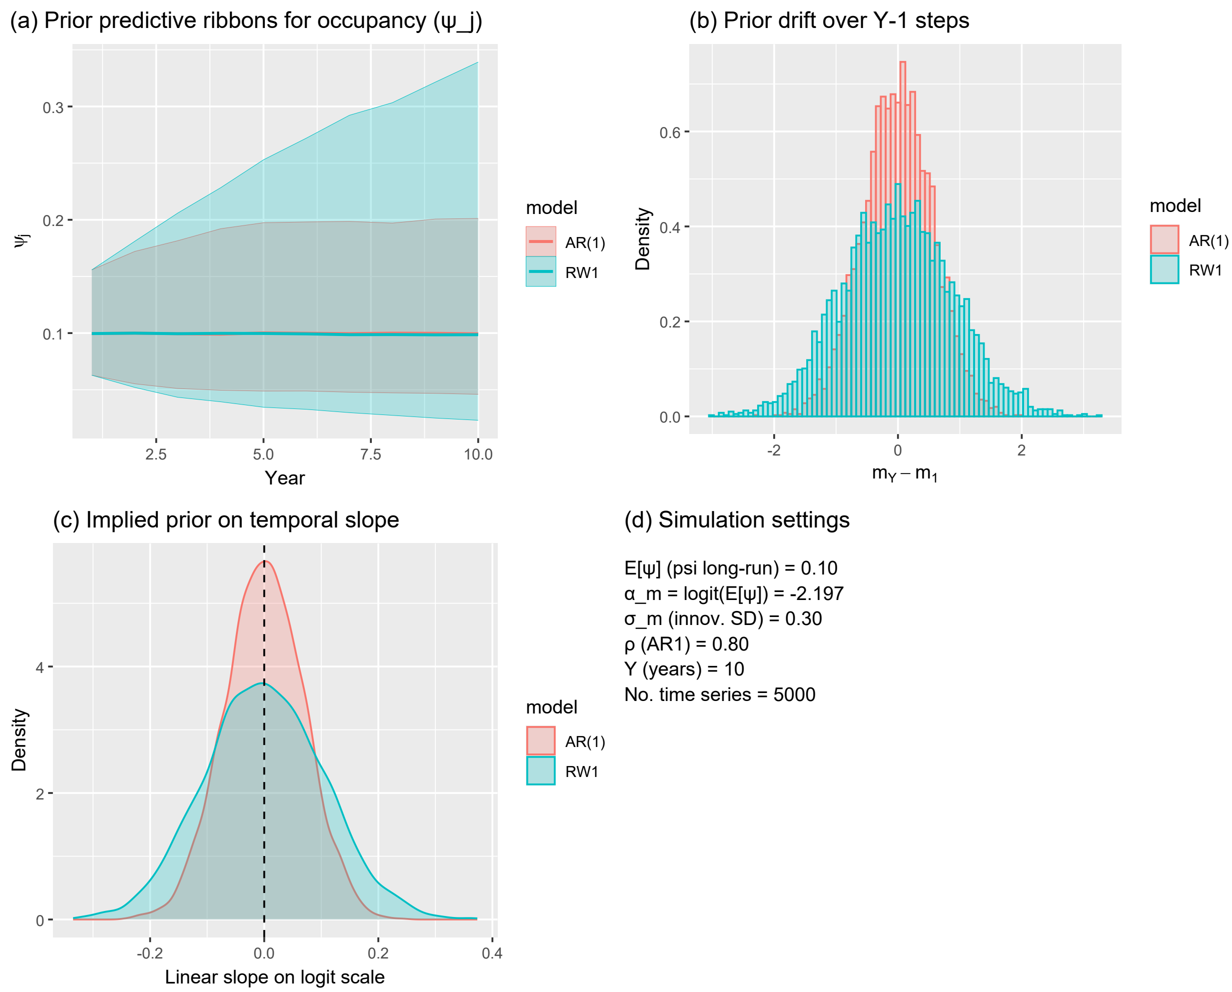

Figure 7: Prior-predictive comparison of temporal priors. (a) Overlaid 5–95% ribbons and medians for ψ_j from 5000 simulated trajectories over 10 years under AR(1) ($\rho = 0.80$) and RW1, with shared baseline $\mathbb{E}[\psi] = 0.10$ and innovation s.d. $\sigma_m = 0.30$ (logit scale). AR(1) bands stabilise around a long-run mean, whereas RW1 widens with time. (b) Distribution of end-to-start difference $m_Y - m_1$ (logit scale): RW1 shows greater drift. (c) Distribution of linear time slopes (logit scale): RW1 induces broader slope variability. (d) Simulation settings.

3.2.2 Other changes

A number of other changes were also trialled during the search for a solution to the issue shown in Figures 5 and 6. Whilst these did not solve the main problem of spurious drift in data-poor species' trends on their own, they were nonetheless judged to be improvements to the stability of the earlier model (v. 4.5; see Appendix 2 for annotated code), and so have been implemented alongside the RW1 to AR(1) shift described above.

3.2.2.1 Plot-year positive-cover draws $\tilde{c}_{i,j}$ (vs a single $\tilde{c}_j$ per year) Although the version of the National Plant Monitoring Scheme model presented in Pescott et al. (2019a) drew estimated latent cover values per plot i and year j, in the model versions used for recent UKBI C7 indicators only a single overall value was drawn per year. This may have led to reduced uncertainty in $\bar{C}_j$, and we have rectified this in version v. 5.0.

3.2.2.2 Truncated half-t standard deviation innovation priors capped at 0.5

Truncating the innovation (i.e. year-to-year standard deviation on the logit scale) limits implausibly large year-to-year jumps under weak data, stabilising trajectories without eliminating the potential for modest changes. A practical trade-off is that extremely abrupt true changes (e.g. >~1 logit) are down-weighted unless the likelihood is very informative.

3.2.2.3 Year 1 occupancy prior centred via Beta(2, 18)

Putting prior mass on low occupancy for the initial state (through $\alpha_m = \text{logit}(\theta),\ \theta \sim \text{Beta}(2, 18)$) regularises very rare species when early years lack information. Abundant species with many detections rapidly overwhelm this prior, so this influence is mainly in data-poor situations.

3.2.2.4 Detection intercept prior Beta(5, 5)

Centring the detection intercept near 0.5 on the probability scale avoids pathological near-0 or near-1 detection baselines when data are sparse or covariates determining detection are missing or unknown. This is expected to improve the calibration of p_a and reduces the posterior sensitivity to weakly identified detection effects.

3.2.2.5 Pareto shape prior for cover precision ϕ_j

A heavier-tailed prior on ϕ_j (shifted from Pareto($k = 1.5$, $c = 0.1$) to Pareto(5, 0.5)) allows both diffuse and sharply peaked Beta distributions for positive cover within a year, adapting to species with very concentrated or very variable cover when present. It regularises precision away from extremes under limited positive-cover data while still permitting high precision if supported.

3.2.2.6 Tighter prior on the cover-detectability slope g_1

We moved from $g_1 \sim \mathcal{N}(0, \tau = 0.1)\ T(0, \infty)$ to $g_1 \sim \mathcal{N}(1, \tau = 0.5)\ T(0, \infty)$. This places more mass on modest, positive effects of cover on detection, shrinking estimates away from extreme near-zero or very large slopes when data are weak, while still allowing the posterior to attenuate towards zero if the data indicate little cover effect.

4 Future work

The current model and all of the changes documented here make the tacit assumption that the locations sampled for any given species-habitat combination are a probability sample of the population about which we wish to make inferences. This assumption is often left unstated in ecological indicator or trend focused work, even though it can seriously undermine inferential goals (Boyd et al., 2022, 2024, 2023). Indeed, reasonable adjustments for time-varying sampling bias have been demonstrated to change species trends, including those on which national indicators are based (Boyd et al., 2025). Given this, the next major piece of work for the NPMS species-habitat trends underlying the UKBI C7 indicator, "Plants of the wider countryside", will be to develop the use of an adjustment model (e.g. Boyd et al., 2024) in order to increase confidence that our trends can be considered representative of the national stock of habitat that they are intended to represent (cf. Pescott et al., 2019b).

5 Appendix 1. Annotated JAGS code (v. 5.0)

```
model{
for (j in 1:Y){ # years
    # re-parameterise Beta distribution for ease of prior setting
    alpha[j] <- mu[j] * phi[j]
    beta[j] <- (1 - mu[j]) * phi[j]
  for (i in 1:N){ # plots
      # draw cover value for each plot/year from year-specific parameters
      # cPos = \tilde{c} above
      cPos[i,j] ~ dbeta(alpha[j], beta[j]) T(1e-5,0.99999)
```

```
      # ZI cover
      # combine new draw of cPos for year j with occupancy
      C[i,j] <- z[i,j] * cPos[i,j]

      # occupancy
      z[i,j] ~ dbern(psi[i,j])
      # year random intercept on occupancy
      logit(psi[i,j]) <- m[j]
      } # end of plots loop
} # end of years loop

## Derived values from state model
for (j in 1:Y){ # number of years
  # annual occupancy
    annOcc[j] <- (sum(z[,j]))/N
  # annual mean cover, including zeros (\bar{C}_j above)
    mC[j] <- mean(C[,j])
  # annual mean occupancy prob
    mPsi[j] <- mean(psi[,j])
}
avgOcc <- mean(annOcc[]) # average annual occupancy

## Plot positive covers
for(k in 1:n.Plot.pos){
  # link positive cover obs to relevant {L_a, U_a}
    cposCens[k] ~ dinterval(cposLatent[k], lims[k,])
  # censored cover informs Beta distribution parameter estimation
    cposLatent[k] ~ dbeta(alpha[year[k]], beta[year[k]])T(1e-5,0.99999)
}

## Observation model for all plot visits ([within-year] detection within plots)
for (a in 1:V2){
    y[a] ~ dbern(py[a]) # detectability influences detection
    py[a] <- z[plotZ[a], yearZ[a]] * p.dec[a] # true state x detectability
    p.dec[a] <- min(max(1e-5, p.Dec[a]), 0.99999) # stop numerical problems

  # detectability model: g1 is cover effect
  # g2 is grazing effect (with missing values)
  # g3 is survey level (no missing values)
    logit(p.Dec[a]) <- g0 + g1*y.cov[a] + g2[x2[a]] + g3[x3[a]]
  # y.cov is c^*_a above
  # y.ind is \delta_a switch above
    y.cov[a] <- yOrig[a]*y.ind[a] + y.int[a]*(1-y.ind[a])
  # transform \hat{c}_a to relevant ordinal category
  # to combine with empirical data
    y.int[a] <- dinterval(y.hat[a], cuts)
  # \hat{c}_a above
    y.hat[a] ~ dbeta(alpha[yearZ[a]],beta[yearZ[a]])T(1e-5,0.99999)
}

### Priors
## Intercept for state model for occupancy
# AR(1) prior for occupancy intercepts m[j]
```

```
# prior for grand mean occupancy
mean.m ~ dbeta(2, 18)T(1e-5,0.99999)
alpha_m <- logit(mean.m)
# AR(1) coefficient
rho_m ~ dunif(0, 1)
# half-Student-t for innovation sd -- +/- 1 logit shifts less likely
sd.m ~ dt(0, 1, 3)T(0,0.5)
tau.m <- 1/pow(sd.m, 2)
m[1] ~ dnorm(alpha_m, tau.m)T(logit(1e-5),logit(0.99999))
for(j in 2:Y){
  m[j] ~ dnorm(alpha_m+rho_m*(m[j-1]-alpha_m), tau.m)T(logit(1e-5),logit(0.99999))
}

## Cover
# AR(1) prior for cover mean on logit-scale (mInt)
phi[1] ~ dpar(5, 0.5) # Pareto prior on precision
mu[1] ~ dbeta(1, 1)T(1e-5,0.99999) # mean prior
alpha_C <- logit(mu[1])
rho_C ~ dunif(0, 1)
# half-Student-t for cover innovation sd
sdC ~ dt(0, 0.5, 3)T(0,0.5)
tauC <- 1/pow(sdC, 2)
mInt[1] <- alpha_C
for(j in 2:Y){
  mInt[j] ~ dnorm(alpha_C+rho_C*(mInt[j-1]-alpha_C), tauC)T(logit(1e-5),logit(0.99999))
  mu[j] <- ilogit(mInt[j])
  phi[j] ~ dpar(5, 0.5)
}

## Detection model
mean.p ~ dbeta(5, 5)T(1e-5,0.99999) # centres p around 0.5
g0 <- logit(mean.p) # transformed
# stronger belief in a positive, moderate-sized cover-detectability effect
# and prevents the model from wandering to very large
# or very small (near-zero) slopes.
g1 ~ dnorm(1, 0.5)T(0,)
# Hierarchical prior for levels of grazing
for (j in 1:g2Levs) { g2[j] ~ dnorm(g2mu, g2tau) }
mean.g2 ~ dbeta(1,1)T(1e-5,0.99999)
g2mu <- logit(mean.g2)
g2tau <- 1/pow(sd.g2, 2)
sd.g2 ~ dt(0, 1, 3)T(0,)

# Hierarchical prior for survey levels
for (j in 1:g3Levs) { g3[j] ~ dnorm(g3mu, g3tau) }
mean.g3 ~ dbeta(1,1)T(1e-5,0.99999)
g3mu <- logit(mean.g3)
g3tau <- 1/pow(sd.g3, 2)
sd.g3 ~ dt(0, 1, 3)T(0,)
} ## END MODEL
```

6 Appendix 2. Annotated JAGS code (v. 4.5)

The annotations below only indicate where changes were made relative to v. 5.0.

```
model{
for (j in 1:Y){
    alpha[j] <- mu[j] * phi[j]
    beta[j] <- (1 - mu[j]) * phi[j]
    cPos[j] ~ dbeta(alpha[j], beta[j]) T(1e-4,0.9999)
  for (i in 1:N){
      C[i,j] <- z[i,j] * cPos[j] ## Change
      z[i,j] ~ dbern(psi[i,j])
      logit(psi[i,j]) <- m[j]
      }
}

for (j in 1:Y){
    annOcc[j] <- (sum(z[,j]))/N
    mC[j] <- mean(C[,j])
    mPsi[j] <- mean(psi[,j])
}
avgOcc <- mean(annOcc[])

for(k in 1:n.Plot.pos){
    cposCens[k] ~ dinterval(cposLatent[k], lims[k,])
    cposLatent[k] ~ dbeta(alpha[year[k]], beta[year[k]])T(1e-4,0.9999)
}

for (a in 1:V2){
    y[a] ~ dbern(py[a])
    py[a] <- z[plotZ[a], yearZ[a]] * p.dec[a]
    p.dec[a] <- min(max(1e-4, p.Dec[a]), 0.9999)
    logit(p.Dec[a]) <- g0 + g1*y.cov[a] + g3[x3[a]] + g2[x2[a]]
    y.cov[a] <- yOrig[a]*y.ind[a] + y.int[a]*(1-y.ind[a])
    y.int[a] <- dinterval(y.hat[a], cuts)
    y.hat[a] ~ dbeta(alpha[yearZ[a]],beta[yearZ[a]])T(1e-5,0.99999)
}

### Priors ###
mean.m ~ dbeta(1,1)T(1e-4,0.9999) ## Change
m[1] ~ dnorm(logit(mean.m), tau0)
tau0 <- 1/pow(sd0, 2)
sd0 ~ dt(0, 1, 3)T(0,) ## Change
for(j in 2:Y){
  m[j] ~ dnorm(m[j-1], tau.m)T(logit(1e-4), logit(0.9999)) ## Change
}
tau.m <- 1/pow(sd.m, 2)
sd.m ~ dt(0, 1, 3)T(0,) ## Change

phi[1] ~ dpar(1.5, 0.1) ## Change
mu[1] ~ dbeta(1, 1)T(1e-4,0.9999)
for (j in 2:Y){
  mu[j] <- ilogit(mInt[j])
  mInt[j] ~ dnorm(logit(mu[j-1]), tauC)T(logit(1e-4), logit(0.9999)) ## Change
```

```
  phi[j] ~ dpar(1.5, 0.1) ## Change
}
tauC <- 1/pow(sdC, 2)
sdC ~ dt(0, 1, 3)T(0,) ## Change

mean.p ~ dbeta(1,1)T(1e-5,0.99999) ## Change
g0 <- logit(mean.p)
g1 ~ dnorm(0, 0.1)T(0,) ## Change

for (j in 1:g2Levs) { g2[j] ~ dnorm(g2mu, g2tau) }
mean.g2 ~ dbeta(1,1)T(1e-5,0.99999)
g2mu <- logit(mean.g2)
g2tau <- 1/pow(sd.g2, 2)
sd.g2 ~ dt(0, 1, 3)T(0,)

for (j in 1:g3Levs) { g3[j] ~ dnorm(g3mu, g3tau) }
mean.g3 ~ dbeta(1,1)T(1e-4,0.9999)
g3mu <- logit(mean.g3)
g3tau <- 1/pow(sd.g3, 2)
sd.g3 ~ dt(0, 1, 3)T(0,)

}
```

7 Acknowledgements

OLP and RJB were supported by the UKCEH National Capability for UK Challenges programme NE/Y006208/1 and by funding from the Joint Nature Conservation Committee to the National Plant Monitoring Scheme, Memorandum of Agreement reference A23-0706-1852.

References

Boyd, R.J., Botham, M., Dennis, E., Fox, R., Harrower, C., Middlebrook, I., Roy, D.B., Pescott, O.L., 2025. Using causal diagrams and superpopulation models to correct geographic biases in biodiversity monitoring data. Methods in Ecology and Evolution 16, 332–344. doi:10.1111/2041-210X.14492.

Boyd, R.J., Powney, G.D., Burns, F., Danet, A., Duchenne, F., Grainger, M.J., Jarvis, S.G., Martin, G., Nilsen, E.B., Porcher, E., Stewart, G.B., Wilson, O.J., Pescott, O.L., 2022. ROBITT: A tool for assessing the risk-of-bias in studies of temporal trends in ecology. Methods in Ecology and Evolution 13, 1497–1507. doi:10.1111/2041-210X.13857.

Boyd, R.J., Powney, G.D., Pescott, O.L., 2023. We need to talk about nonprobability samples. Trends in Ecology & Evolution doi:10.1016/j.tree.2023.01.001.

Boyd, R.J., Stewart, G.B., Pescott, O.L., 2024. Descriptive inference using large, unrepresentative nonprobability samples: An introduction for ecologists. Ecology 105, e4214. doi:10.1002/ecy.4214.

Pescott, O., Walker, K.J., Powney, G., 2019a. Developing a Bayesian Species Occupancy/Abundance Indicator for the UK National Plant Monitoring Scheme. Unpublished Report to JNCC/Defra.. NERC/Centre for Ecology & Hydrology and BSBI. Wallingford, UK.

Pescott, O.L., Jitlal, M., Smart, S.M., Walker, K.J., Roy, D.B., Freeman, S.N., 2016. A comparison of models for interval-censored plant cover data, with applications to monitoring schemes. PeerJ Preprints 4, e2532v1. doi:10.7287/peerj.preprints.2532v1.

Pescott, O.L., Walker, K.J., Harris, F., New, H., Cheffings, C.M., Newton, N., Jitlal, M., Redhead, J., Smart, S.M., Roy, D.B., 2019b. The design, launch and assessment of a new volunteer-based plant monitoring scheme for the United Kingdom. PLoS ONE 14, e0215891. doi:10.1371/journal.pone.0215891.

Plummer, M., 2017. JAGS Version 4.3.0 User Manual.