Ecological Modelling 500 (2025) 110948

Contents lists available at ScienceDirect

Ecological Modelling

journal homepage: www.elsevier.com/locate/ecolmodel

Calibrating multi-constraint ensemble ecosystem models using genetic algorithms and Approximate Bayesian Computation: A case study of rewilding at the Knepp Estate, UK

Emily Neil [*], Ernesto Carrella, Richard Bailey

School of Geography and the Environment, University of Oxford, United Kingdom

ARTICLE INFO

Keywords:
Ensemble ecosystem modelling
Lotka-Volterra
Rewilding
Species reintroductions
Calibration
Genetic algorithms
Optimisation

ABSTRACT

This paper presents a new ensemble ecosystem model (EEM) which predicts the impacts of species reintroductions and optimises potential future management interventions at the Knepp Estate rewilding project, UK. Compared to other EEMs, Knepp has a relatively high level of data availability that can be used to constrain the model, including time-series abundance data and expert knowledge. This could improve the realism of outputs and enable more nuanced and context-specific management intervention recommendations. Calibrating EEMs can be challenging, however, and as the number of constraints increases, so does the complexity of the model fitting process. We use a new Genetic Algorithm-Approximate Bayesian Computation (GA-ABC) approach wherein GA outputs are used to inform the prior distributions for ABC. To reduce the parameter search space, we fixed twelve parameters - the consumer self-interaction strengths $\alpha_{i,i}$ and negative growth rates – based on theoretical assumptions. While the GA-ABC method proved effective at efficiently searching the parameter space and optimising multiple constraints, it was computationally intensive and struggled to identify a broad range of outputs. Ultimately, this led to an ensemble of models with similar trajectories. Several potential ways to address this are discussed. Our results reinforce the findings of previous studies that the EEM methodology has potential for guiding conservation management and decision-making. Outputs suggest that reintroducing large herbivores was key to maintaining a diverse grassland-scrubland-woodland ecosystem, and optimisation experiments informed species characteristics and stocking densities needed to achieve specific goals. Ultimately, refining the EEM methodology to improve calibration and facilitate the integration of additional data will enhance its utility for ecosystem management, helping to achieve more effective and informed outcomes.

1. Introduction

Rewilding is an innovative approach to ecosystem restoration, aiming to re-establish natural processes and food web interactions - often through species reintroductions - to restore lost ecological functions and create resilient, self-sustaining ecosystems (Carver et al., 2021). A growing body of evidence supports its benefits (Navarro and Pereira, 2012; Svenning, 2020), but concerns remain due to limited empirical research and the challenges of predicting its outcomes, including possible ecological consequences (Nogués-Bravo et al., 2016; Svenning et al., 2016). Ecosystems are usually complex, and perturbations can have widespread and unexpected knock-on effects, potentially producing unintended consequences. Species reintroductions, for example, may result in declines or extirpations of other species (Lovari et al., 2014; Pace et al., 1999), or significant shifts in species composition (Johnson and Cushman, 2007). Such impacts are often dependant on management decisions, such as the functional groups or population sizes of reintroduced species (Newton et al., 2013; Scheffer et al., 2001;Tuomi et al., 2019). With rewilding projects on the rise (Jepson, 2019), there is a growing need for analytical approaches that can forecast the outcomes of different management interventions.

A key challenge in predicting ecological outcomes is the limited availability of long-term data on species' abundances and interactions. In many cases, the only available information comes from expert knowledge or sporadic, short-term surveys, introducing significant uncertainty when trying to model complex ecosystem dynamics. Ensemble ecosystem models (EEMs) (Baker et al., 2017) are well-suited for addressing these data limitations. The methodology is conceptually

---

[*] Corresponding author.
E-mail address: emily.neil@worc.ox.ac.uk (E. Neil).

https://doi.org/10.1016/j.ecolmodel.2024.110948
Received 6 May 2024; Received in revised form 10 November 2024; Accepted 11 November 2024

similar to Approximate Bayesian Computation (ABC) with rejection sampling (Pritchard et al., 1999): a large ensemble of outputs is generated by randomly sampling parameter values from prior distributions, retaining only those that produce plausible ecosystems. This provides a broad range of potential ecosystem dynamics, offering valuable insights into the potential impacts of various management actions made under uncertainty. EEMs have been applied to a variety of ecological management problems, including species reintroductions (Baker et al., 2017, 2019; Bode et al., 2017), invasive species management (Rendall et al., 2021), and assisted migrations (Peterson and Bode, 2021).

EEMs typically use constraints to exclude models that do not result in stable and feasible ecosystems. A stable ecosystem can recover from small perturbations, while a feasible one ensures all species maintain positive abundances at equilibrium. While these have been useful for calibration in the absence of data, questions have arisen over the validity of these assumptions for real-world ecosystems (Vollert et al., 2024a). The Knepp Estate rewilding project in the UK offers an opportunity to instead constrain model outputs with the integration of relatively comprehensive datasets, including both expert knowledge and over a decade of time-series abundance data, which can indirectly inform the strength of species interactions (Baker et al., 2019; Raymond et al., 2011). This could enhance the realism of outputs and enable more nuanced and context-specific management intervention recommendations. Still, uncertainties and gaps remain with the available data at Knepp, making EEMs an appropriate choice for navigating the complexities of this relatively "data-rich" but still technically data-limited system.

However, calibrating EEMs is often a significant challenge (see Vollert et al., 2024b, Peterson and Bode, 2021), and as the number of constraints increases, so does the complexity of calibration due to the curse of dimensionality (Bellman, 1961). The standard EEM approach is to randomly sample the parameter space until an ensemble of stable and feasible ecosystems is generated (Baker et al., 2017), but this can be inefficient: Peterson & Bode (2021), for example, found that in a network of 15 species, there was less than a 1 in 1 million probability of randomly sampling a feasible and stable ecosystem. Baker et al. (2019) developed an EEM that could integrate constraints based on time-series abundance data, and found that uninformed choices for parameter distributions resulted in no outputs which could accurately reproduce the complex observed dynamics. The authors overcame this by using an iterative search process. This involved widening the data-derived constraints until 100 acceptable parameter sets could be identified, fitting distributions to the outputs to inform future searches. The constraints were then gradually tightened until a set of 100 parameters were identified which could pass the original constraints. This approach led to many very similar parameter sets, so to increase variability, the authors added random noise to the distributions at each iteration. While effective, this approach can be computationally and labour intensive, as it requires multiple rounds of adjustment and refinement to identify suitable parameter sets. Metaheuristic algorithms, such as genetic algorithms (GAs), could offer an alternative approach and potentially speed up the calibration process by efficiently searching large parameter spaces while navigating any constraints (Abdel-Basset et al., 2018).

In this paper, we extend the work of Baker et al. (2019) by applying EEMs to the Knepp Estate, replacing stability and feasibility assumptions with 33 constraints derived from time-series abundance data, and 4 additional 'reality check' constraints based on expert knowledge and literature from similar ecosystems. To efficiently calibrate the model within this large parameter space, we use a new Genetic Algorithm-Approximate Bayesian Computation (GA-ABC) approach, in which GA outputs are used to inform prior distributions for ABC. GAs are search methods based on the principle of evolution via natural selection: the 'fittest' candidates - i.e., those parameter sets that closely fit the data - are selected to produce the next generation. Offspring have 'genes' which are similar to the parents' but can be subjected to mutation. Over successive generations, the community evolves towards an optimal solution (Bremermann, 1958; Fraser, 1957; Holland, 1975). They have been applied to a wide range of optimisation problems (Ross and Corne, 1994), including ecological ones (Hamblin, 2013; Ruxton and Beauchamp, 2008).

GAs provide many advantages for EEM calibration: by evolving solutions based on fitness evaluations, they search the parameter space more efficiently than both the standard random sampling approach used in EEMs (Baker et al., 2017) and potentially also the iterative process used by Baker et al. (2019). GAs can optimise many constraints simultaneously and are well-suited to nonlinear problems, like those commonly observed in ecosystems, where dynamics can be complex. Sampling is global rather than local, reducing the risk of becoming trapped in local minima (Massart et al., 1998) and avoiding reliance on an assumed starting model (Gallagher and Sambridge, 1994). This is particularly valuable for complex ecosystem models with large parameter spaces and many parameter uncertainties. Additionally, GAs can be run multiple times to generate a diverse ensemble of solutions, offering a broader exploration of potential parameter sets and avoiding convergence on many similar models. ABC can further improve outputs by identifying more acceptable models and offering probabilistic insights into uncertainty (Pritchard et al., 1999).

Our study represents a methodological advancement by integrating a relatively high level of real-world data into a new EEM of the Knepp Estate, and by providing a new solution for the calibration of models with complex constraints from multiple sources. Our objectives are twofold: (1) to evaluate the performance of EEMs and the GA-ABC approach in a case study with a relatively high amount of data availability from multiple sources, and (2) to provide practical guidance for others using this methodology in similarly data-rich systems. Additionally, we discuss the broader utility of EEMs for rewilding practice, exploring their potential to provide new insights and optimise management interventions.

2. Methods

2.1. Rewilding at the Knepp Estate

Knepp is a 3500 acre estate in West Sussex, England (Tree, 2018). Once a high-intensity arable and dairy farm, in 2001 the land was devoted to a pioneering project inspired by rewilding principles and the wood-pasture hypothesis (Vera, 2000). The wood-pasture hypothesis proposes that the primeval European landscape would have been wood pasture, characterised by a shifting mosaic of grassland, scrubland and woodland, with dynamics driven by the grazing, browsing, and trampling of large herbivores (Harding and Rose, 1986; Vera, 2000). Spatial association with unpalatable plants (e.g., thorny scrub) enables palatable plants (e.g., tree saplings) to regenerate in grasslands maintained by large grazers. The palatable plants then out-shade the species that facilitated them, leading to cyclical turnover (Olff et al., 1999). Many rewilding initiatives, including Knepp, are introducing large herbivores to maintain these dynamic and biodiverse wood pasture habitats (Smit et al., 2015).

Knepp is divided into three blocks: Northern, Middle, and Southern, each of which varies in management strategy, herbivore assemblage and density, and the year in which rewilding began. This paper focuses specifically on the 4.45 square-kilometre Southern Block, which has the most time-series data. In 2009, Exmoor ponies, Old English longhorn cattle, and Tamworth pigs were reintroduced to act as domesticated proxy species for the tarpan, aurochs, and wild boar, respectively. Fallow deer were then reintroduced in 2010, and red deer in 2013. Roe deer were already naturally occurring at the site in low numbers. Reintroduced species are free roaming within the blocks of the estate and maintained at low densities relative to conventional farms and conservation grazing (Barkham, 2018; Natural England, 2018). This is done via culling (harvesting for meat) or the removal of breeding males, to allow for vegetation regeneration. There are no large carnivores present at Knepp, and the culling regime is primarily guided by: 1) animal welfare issues, i.e., reducing the population to a level that can survive long, harsh winters given fodder availability, and 2) a desire for a wood-pasture habitat (Tree, 2018).

Since the project began, there have been significant changes to habitat composition in the Southern Block, with shifts from arable fields and grassland to a complex mosaic of predominantly scrub and woodland (Schulte to Bühne et al., 2022). This increase in vegetation cover has likely had impacts on ecosystem services such as carbon sequestration and flood mitigation, and facilitated increases in invertebrate, bird, and small mammal diversity, including the return of rare species such as nightingales (*Luscinia megarhynchos*), turtle doves (*Streptopelia turtur*), and purple emperor butterflies (*Apatura iris*) (Kurwald, 2017; Tree, 2018).

Knepp has time-series data available from various sources: a baseline survey of species and habitat types at Knepp was undertaken in 2005 before species reintroductions took place (Greenaway, 2006), and data on each reintroduced species' populations are available on a yearly basis between 2009 and 2015, and monthly between 2015 and 2021. Data are also available on habitat composition from various studies conducted between 2001 and 2020. An ecologist is on site to provide expert knowledge.

2.2. *The ecosystem network*

The model, developed in Python and openly available on GitHub (Neil et al., 2024), simulates the population dynamics of interacting species at Knepp. It is based on a network of consumers and habitat components which was informed by the wood-pasture hypothesis, quantitative data from Knepp and similar sites, and qualitative expert opinion from the resident ecologist (Kirby, 2004; Olff et al., 1999; Putman, 1986; Penny Green, *pers comm.* 2020). There are seven nodes representing consumers: Old English longhorn cattle, Exmoor ponies, fallow deer, red deer, roe deer, Tamworth pigs, and a hypothetical new species whose impacts will be explored by the model; and three habitat components: grassland, woodland, and thorny scrub.

Interactions can be trophic (predator-prey, or in the context of Knepp, herbivory) or non-trophic (e.g., competition, facilitation) (Fig. 1, Appendix A). Consumers are assumed to negatively impact, and be positively impacted by, habitat types via herbivory. Larger habitat types are assumed to negatively impact smaller types via competition for light, for example woodland is assumed to outcompete grassland and thorny scrub (Olff et al., 1999). As per the wood-pasture hypothesis, thorny scrub is assumed to positively impact woodland by protecting saplings from herbivory and facilitating regeneration (Smit et al., 2006; Smit and Ruifrok, 2011; Vera, 2000). Although other elements could be important for the system, such as the soil microbiome and nutrient cycling, they are not added primarily since they are outside the focus of the wood-pasture hypothesis, but also due to lack of data.

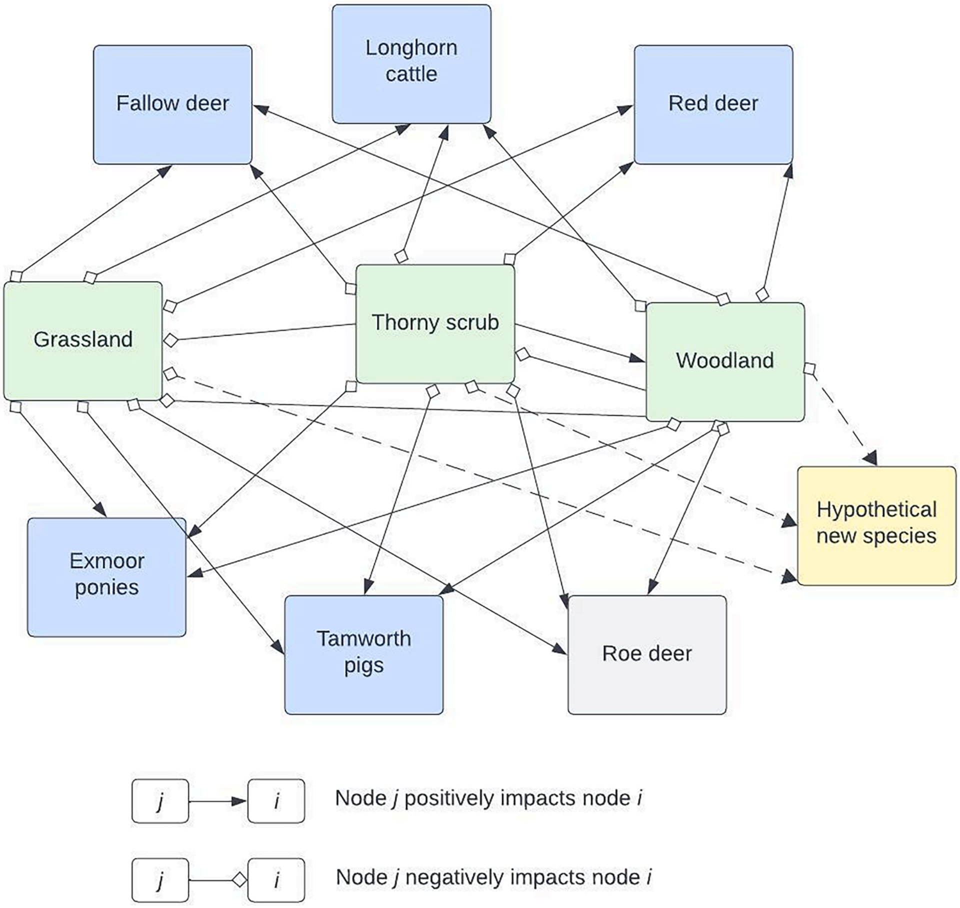

Fig. 1. An ecological network of the Knepp Estate. Solid lines are the interactions between ecosystem elements which are currently present in the system. Blue boxes are consumers that were reintroduced to Knepp and whose numbers are continually managed via culling. Grey boxes are consumers that were already present at Knepp and whose populations are un-managed. Consumers are assumed to negatively impact, and be positively impacted by, each of the three habitat types (green boxes). Habitats can negatively impact each other via competition for light. Thorny scrub is assumed to positively impact woodland via facilitation of saplings. The model also explores potential impacts of a hypothetical new species (yellow box) which is assumed to have similar interactions (dashed lines) to the other consumers.

2.3. Model construction

Fig. 2 outlines a flowchart of methods used for model construction and calibration. The model is based on the generalised Lotka-Volterra equations (Eq. (1); Lotka, 1925; Volterra, 1926):

$$\frac{dn_i}{dt} = r_i n_i + \sum_{j=1}^{N} \alpha_{i,j} n_i n_j \qquad \text{(Eqn. 1)}$$

Where n_i is the total amount of ecosystem element i (relative population size, i.e., relative number of individuals or percentage habitat type), r_i is the growth rate (growth of the relative population per year), $\alpha_{i,j}$ represents the effect of the interaction between nodes i and j on the node i (relative number of individuals^{-1}year^{-1}), and N is the total number of ecosystem elements. Initial population sizes for each node (e.g., the initial percentages of each habitat type or number of roe deer in 2005, and the initial number of each consumer upon its reintroduction) are rescaled to 1, facilitating comparative analysis across different species by focusing on relative rather than absolute changes. Initial population sizes before rescaling to 1 are outlined in Appendix A.

Habitat components have positive r_i values between 0 and 1 and are assumed to be self-limiting, with self-interaction strengths $\alpha_{i,i}$ between -1 and 0 (Baker et al., 2017, 2019; Bode et al., 2017). Consumer species are assumed to have negative r_i values and self-interaction strengths $\alpha_{i,i}$ equal to zero (Fig. 2b). This differs from other EEM models in the literature, which generally have positive intrinsic growth rates and negative self-interaction strengths $\alpha_{i,i}$ for all ecosystem elements. The reason for this difference is that, in this model, it is important for consumer populations to collapse in the absence of habitat components. With positive intrinsic growth rates, consumer populations can survive without any inputs (food sources) for several years; while this is a standard assumption made in many EEM papers, it was considered too unrealistic for this case. Interaction strengths are assumed to be between -1 to 5, depending on the sign of the interaction. EEMs typically have interaction strengths of between -1 and 1, but in our case, we set the upper bound at 5 to allow these interactions to sufficiently counterbalance consumer negative growth rates and ensure population persistence.

The models cover three time-periods: the years before large herbivores were reintroduced (2005–2009), the years after large herbivores were reintroduced (2009–2020), and future projections (2021–2120). Splitting the model in this way allows for the assessment of what may have occurred if species reintroductions had not taken place.

2.4. Calibration using genetic algorithms and Approximate Bayesian Computation

The aim is to generate an ensemble of parameter sets that replicate the empirical data while also passing several simple ‘reality check’ constraints to ensure the model is fundamentally theoretically sound. Time-series data on each reintroduced species’ abundance are available from Knepp on a yearly basis between 2009 and 2015, and monthly between 2015 and 2020. At the start of each simulated year between 2009 and 2020 (post-species reintroductions), the average annual population value for each reintroduced species is forced into the model to simulate culling. To ensure their populations grow at realistic rates, each simulation from 2015 onwards (when monthly consumer abundance data is available) passes through constraints informed by the maximum population values per year. Two time-series abundance data points are available to inform constraints for each habitat component and for roe deer. Further details on the time-series data, constraints, and forcings are provided in Appendix A.

In addition to the data-derived constraints, four reality checks are explored (outlined below). Constraint 1 is passed by setting fixed parameters. Constraints 2–4 rely on additional filters, as they are not parameter-specific but instead dependant on model outputs and interactions between the parameters (i.e., they could not be accommodated by changing the priors alone).

1. Constraint 1: If all habitat components are removed, consumer populations should collapse towards ≈0 within one year, reflecting the simple physiological needs of the animals. We define “collapse" to mean falling to an arbitrary small value of 0.01. To enforce this, the intrinsic growth rate r_i for consumers is dynamically updated to ensure this condition is met as populations change. Assuming consumers have self-interaction strengths of zero and no additional interactions, and $n_i(t-1)$ is non-zero, solving for a decay to 0.01 within 1 year yields the equation:

$$r_i(t) = log\left(\frac{1}{100 * n_i(t-1)}\right)$$

2. Constraint 2: If the model is run for 100 years with no consumers, thorny scrubland should increase at the onset, but then eventually be outcompeted by woodland. This is based on expert knowledge and empirical evidence from other rewilding projects in England which found that agricultural fields can become shrubland-dominant after approximately 15 years, and closed-canopy woodland after 40–50 years (Broughton et al., 2021). The constraint gives some flexibility around these time-frames and assumes thorny scrub should be greater than woodland and grassland at year 15, and woodland should be greater than thorny scrub and grassland at year 100.
3. Constraint 3: If many herbivores are artificially maintained within the system for a long period of time, all habitat types should decline

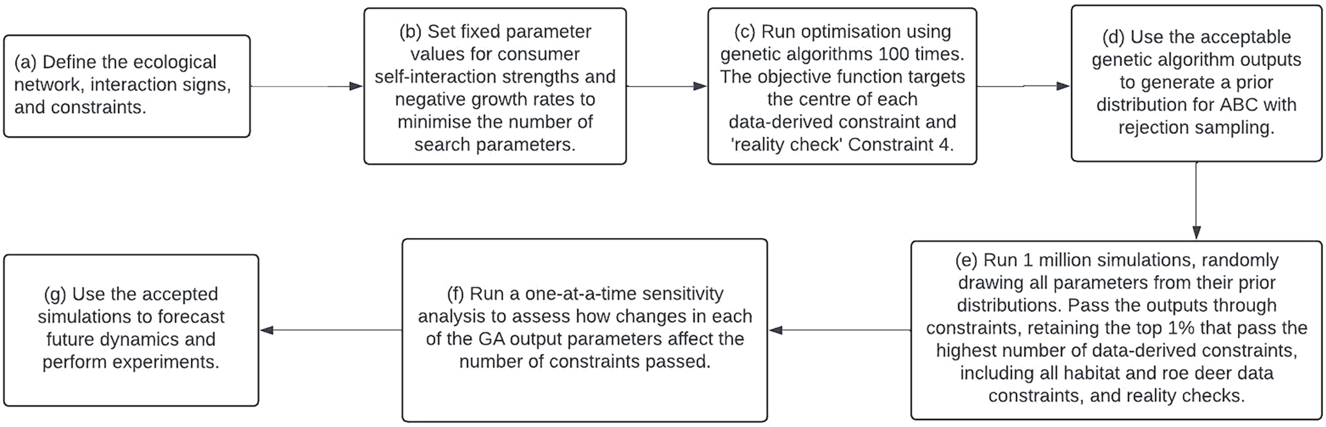

Fig. 2. Flow diagram for the procedures used to model the ecosystem dynamics at the Knepp Estate, specifically model construction (boxes a-b), calibration using genetic algorithms and Approximate Bayesian Computation (boxes c-e), and experiments (boxes f-g).

to < 0.01 within 100 years. This is based on studies in the UK which show excessive grazing can reduce structural complexity and species richness, and suppress tree regeneration (Anderson and Radford, 1994; Kinnaird, 1974; Pollock et al., 2005). To simulate this constraint, reintroduced species are added at twice their 2020 stocking densities. Consumers' interaction strengths $\alpha_{i,j}$ with grassland are set to the maximum of 5 and r values set to 0. This ensures that consumer populations grow or stay constant, since the model also assumes no negative interactions between consumers and no self-limitation for consumers. This could theoretically represent a scenario in which consumers are given supplemental feeding for the purpose of keeping them in the area.

4. Constraint 4: In the absence of consumers or habitat competition, grassland, scrubland, and woodland should each level out near the equivalent of 100 % within 100 years (values of 80–120 % at year 100 were considered acceptable), bearing in mind that initial conditions are rescaled to 1 in the model. It is theoretically possible that model output values could be far above or below 100 % without this constraint.

The inclusion of these constraints makes filtering too demanding to pass a simple accept/reject scheme. Therefore, optimisation using genetic algorithms (GAs) is used to explore the space (46 parameters) and identify a parameter set that produces outputs closely matching the empirical data from Knepp while also passing the reality checks (Fig. 2c). The objective function (Eq. (2)) focuses on minimising the value a, which is the deviation between the model-predicted outputs (x_i) and the target values derived from the empirical data (y_i), where K is the total number of constraints. This approach targets the centre of each data-derived constraint (y_i): for example, the woodland abundance datapoint in 2020 is 21.5 %, which, when the initial condition of 5.8 % is rescaled to 1, corresponds to a value of 3.7 (Appendix A). In the ABC process, acceptable outcomes are assumed to fall within ±5 of 21.5 %, with a rescaled range of 2.8–4.6. In the GA process, the objective function targets the middle of this range (3.7), ensuring the models match the empirical data as closely as possible.

In addition to the abundance data constraints, the objective function targets 'reality check' constraint 4, which focuses on limiting maximum values for nodes rather than more general ecosystem dynamics and is therefore assumed to be the most challenging to satisfy. For habitat nodes governed by this constraint, the objective function targets the middle of the acceptable 80–120 % range (centred at 100 %). Since the initial conditions are rescaled to 1, the target values corresponding to 100 % for each habitat are: 1.1 for grassland, 23.3 for scrubland, and 17.2 for woodland.

$$a = \sum_{i=1}^{K} \left(\frac{(x_i - y_i)}{y_i} \right)^2 \tag{Eqn 2}$$

One hundred separate GAs are run. Outputs which pass all reality check constraints, as well as all habitat and roe deer abundance constraints, are used to construct a prior distribution for ABC with rejection sampling (Pritchard et al., 1999) to identify additional suitable parameters and model outputs (Fig. 2d). The second requirement exists because reintroduced consumer species are continuously observed and have more data constraints (4–6 per reintroduced species, versus 2 for roe deer and each habitat type), so the outputs may skew towards correctly fitting these. However, a key aspect of the model is to understand the effect of grazing animals on habitat composition, and a usable model must therefore be able to correctly replicate habitat dynamics. 1 million simulations are then run by randomly drawing all parameter values from these distributions, and the top 1 % of simulations that pass the highest number of data-derived constraints, including all habitat and roe deer data constraints and reality checks, are retained (Fig. 2e).

The posterior joint distribution of parameters is built using only the accepted simulations. Accepted simulations are used to run sensitivity analyses (Fig. 2f), forecast the range of potential future dynamics of the ecosystem in the long-term, and to explore questions (Fig. 2g) such as: how have species reintroductions impacted the ecosystem, and how might they continue to impact it in the future? What characteristics would a new species need to have, or what changes in current stocking densities would need to occur, to engineer the system towards a particular end goal?

3. Results

3.1. Accepted simulations

One hundred separate GAs were run to identify a range of parameter sets (Fig. 2c). The outputs exhibited relatively good fits to the data, however, none were identified that passed all constraints (Appendix B). Constraint 2 and habitat abundances were the most difficult to pass, with zero accepted runs.

Numerical optimisation becomes more efficient, particularly when dealing with a large number of parameters, by manually narrowing the range of possible parameter values before optimisation begins. This is done here as a next step, by removing, before feeding to the optimiser, large parts of the parameter space which would obviously result in unrealistic population dynamics. Given enough time and computational resources, one would expect a GA (or any other optimisation algorithm) to discover this property of the parameter space independently, but by manually removing nonsensical options from the parameter space (nonsensical given our understanding of the ecosystem which GAs lack) we can greatly speed up the calibration process. To achieve the desired system behaviour, we used trial and error to adjust certain parameters:

1. Woodland r is assumed to be ≤ 0.15, and the positive interaction with thorny scrub ≤ 0.01, to prevent it from becoming 100 % of the system in < 45 years; this would be considered unrealistically high growth (Broughton et al., 2021). The negative self-interaction strength $\alpha_{i,i}$ is assumed to be -0.01 to -0.005 to allow woodland to level out near 100 % in the absence of herbivory (Constraint 4). Since the initial condition of 5.8 % was rescaled to 1, a model output of 17.2 would be equivalent to 100 % woodland.
2. Thorny scrub r is assumed to be ≤ 0.5 to prevent it from becoming 100 % of the system in <15 years; this would be considered unrealistically high growth (Broughton et al., 2021). The negative self-interaction strength $\alpha_{i,i}$ is assumed to be -0.025 to -0.01 to allow scrubland to level out near 100 % in the absence of habitat competition and herbivory (Constraint 4). Since the initial condition of 4.3 % was rescaled to 1, an output of 23.3 would represent 100 % scrub.
3. Grassland is assumed to have high growth rates (≥ 0.75) so it can increase towards 100 % in a year in the absence of habitat competition and herbivory. The negative self-interaction strength $\alpha_{i,i}$ is bounded between -1 to -0.75 to allow grassland to level out near 100 % in the absence of habitat competition and herbivory (Constraint 4). Grassland had an initial condition of 89.9 % which was rescaled to 1. An output value of 1.1 would be equivalent to approximately 100 %.
4. To prevent grassland, thorny scrub, and woodland from falling towards zero by the 2020 data constraints, grassland and scrub are assumed to be impacted $\geq$-0.05 by herbivory and habitat competition-related constraints, and woodland $\geq$-0.025 by herbivory (woodland is assumed to not be affected by habitat competition).
5. Consumers are assumed to benefit 0–0.5 from woodland and thorny scrubland to prevent their populations from experiencing unrealistically rapid growth as these habitat nodes increase. To prevent their populations from collapsing in grassland-dominant environments, consumers are assumed to benefit $\geq$2.5 from grassland.

After manual adjustments and an additional 100 separate GA runs (repeating step c as shown in Fig. 2), one parameter set was identified which satisfies constraints 1–4 and all habitat and roe deer abundance constraints, and optimises Eq. (2), passing 94 % of data constraints (Appendix B). Ideally, the ABC prior would be informed by all 100 GA outputs. However, since only 1 passed all constraints, the neighbourhood of this parameter set is explored to find additional suitable parameters and model outputs. This is done so by assuming, for each parameter, a uniform prior ranging ±10 % of the value returned by the GA (Fig. 2d). Of the 1 million runs, 637,020 passed the first set of data constraints in 2009, and 204,509 passed the reality checks and final roe deer and habitat constraints in 2020. Taking the top 1 % left a total of 2045 accepted simulations (Fig. 2e; Fig. 3). Each data-derived constraint was passed by at least 15.1 % of the accepted simulations (Appendix C). The two most difficult to pass were fallow deer populations in 2018 (16.8 %), and 2019 (15.1 %).

A one-at-a-time sensitivity analysis (ten Broeke et al., 2016) was undertaken to assess how changes in parameters affect the number of abundance constraints passed (Fig. 2f). Each parameter in the accepted GA output was varied ±50, ±10, ±5, or ±1 % around its value, the model run with the new value, and the percentage of passed constraints recorded. Changes in eight parameters: all three habitat growth rates, grassland and scrubland negative self-interactions, and three interaction parameters (woodland impact on ponies, grassland impact on fallow

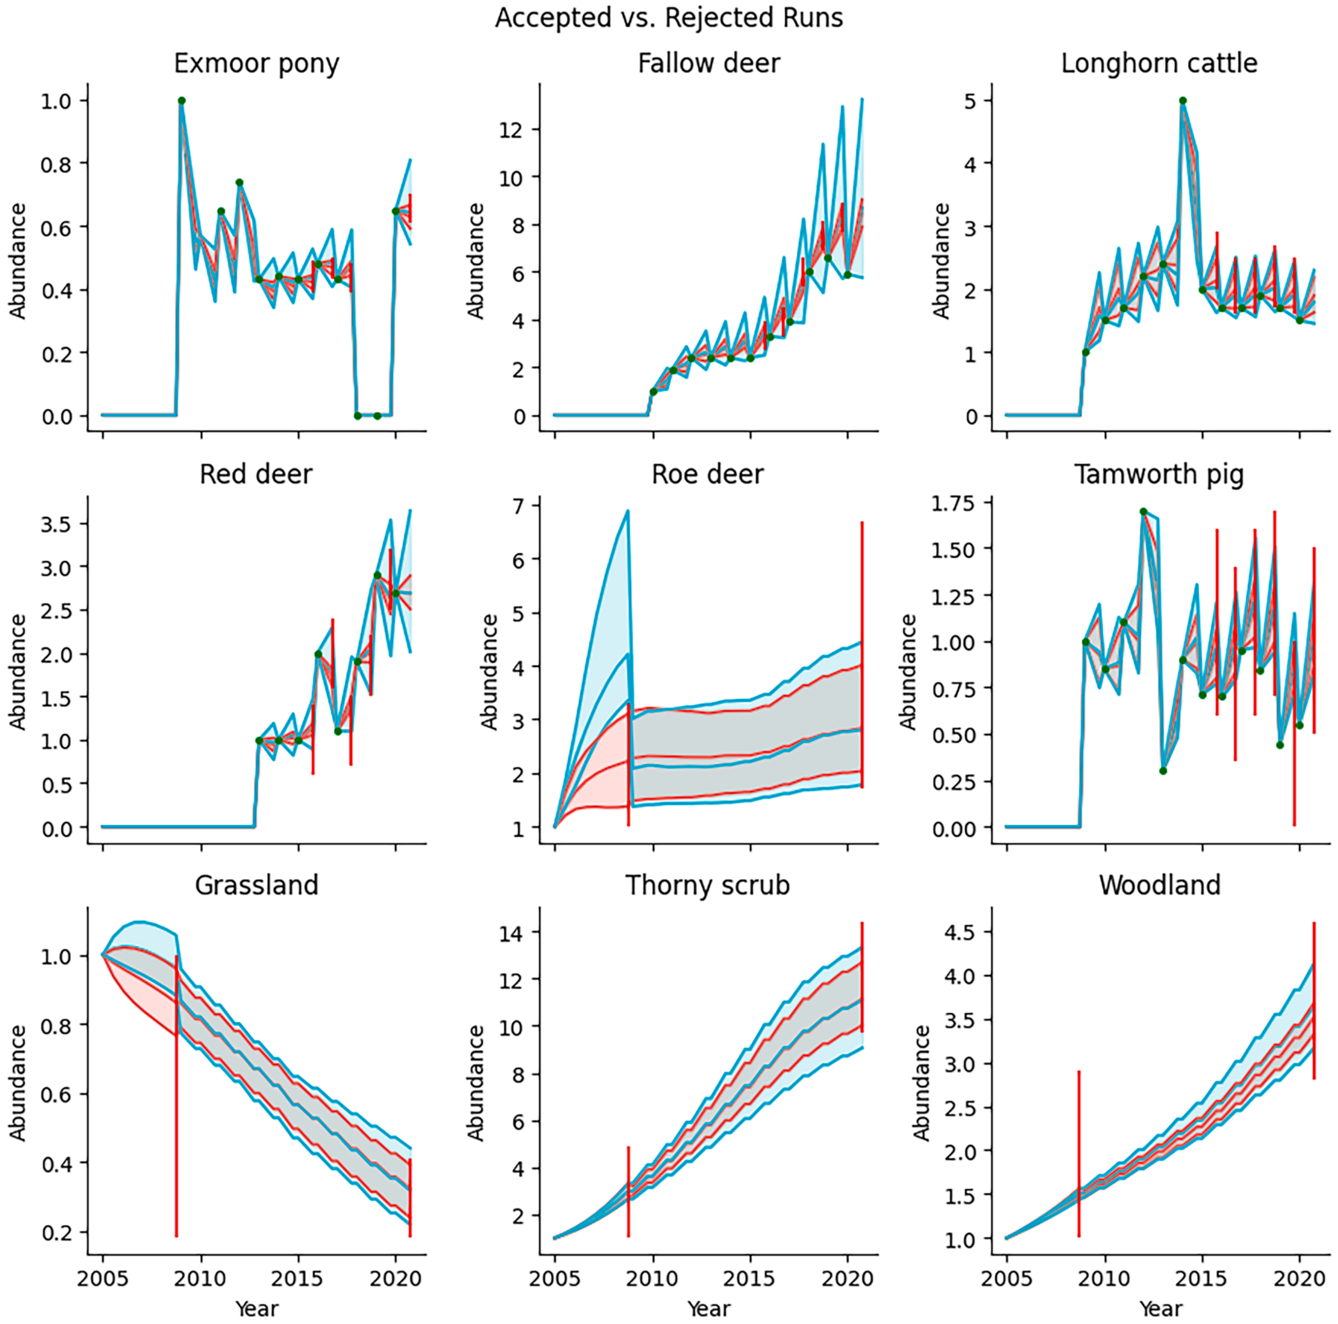

Fig. 3. Time-series of accepted (red, representing 2045 runs) vs. rejected (blue, representing 997,955 runs) runs for each ecosystem element in the pre- and post-reintroduction simulations. Each graph depicts the median and 95th percentiles (shaded areas) of the model outcomes. Red vertical lines represent the data constraints each run passed. Dots indicate external forcings applied to the models (e.g., the number of consumers culled at each time point), which affect only the reintroduced species (Exmoor pony, fallow deer, Longhorn cattle, red deer, and Tamworth pigs). These forcings account for the 'jumpy' appearance of the abundance trajectories for these species. This figure tracks relative change over time as all initial conditions are rescaled to 1 (see Appendix A for the original abundance data values).

deer, and scrubland impact on grassland) had the highest impact on the percentage of passed constraints (Fig. 4, Appendix C). Future data collection could be directed to these parameters to gain the most benefit per observation (LeBauer et al., 2013).

3.2. Simulation experiments

The accepted simulations are used to replicate the ecosystem dynamics from 2005 to 2020 and project 100 years into the future, and to assess what may have occurred if species reintroductions had not taken place (Fig. 2g). At the beginning of each year when forecasting, reintroduced species' populations are forced to their 2020 populations, as Knepp seeks to maintain this stocking density in future (Penny Green, *pers comm.* 2020). Results suggest that with current stocking densities, the system will become primarily (60 %) woodland in the next 50 years, with scrubland being maintained at about 25 % and grassland at approximately 10 % (Fig. 5). Had species reintroductions not occurred, the ecosystem would have become almost exclusively (90 %) woodland despite the presence of roe deer, with scrubland and grassland each declining to <5 %. Roe deer numbers are about twice as high if reintroduced species are not present compared to current dynamics.

Next, two different groups of simulations are run to assess how management can engineer the ecosystem towards a given end goal. First, the model is run 100 years into the future and stocking densities of all five reintroduced species are varied simultaneously from 0 to 1000 % around their current levels. Second, a hypothetical new species is reintroduced. The new species is assumed to have similar interactions to the other consumers: it negatively impacts all three habitats but does not directly interact with the other consumers present in the system. Its parameters regarding the impacts it has on grassland, scrubland, and woodland (with bounds of -1 to 0), and the amount it gains from grassland, scrubland, and woodland (bounds of 0 to 5) are varied. It is assumed that the new species' populations are controlled as per the other reintroduced species, so abundances are forced back to the initial condition of 1 at the beginning of each year. The hypothetical goals for both experiments are reducing woodland regeneration from 60 % in 100 years to target values of either 50 %, 45 %, or 25 %. GAs are re-run for these targets, with the objective function as defined in Eq. (2).

Results suggest that all stocking densities would need to increase to reduce woodland regeneration (Fig. 6). Outputs have increases of 18–174 % compared to current stocking densities when engineering towards 50 or 45 % woodland, with larger increases of 220–455 % required when targeting 25 % woodland.

Introducing a new species to reduce woodland regeneration showed significant impacts over the 100-year projections. Optimiser outputs suggest this hypothetical new species often has higher impacts on habitats than the other consumers: negative interaction strengths between the new species and grassland are up to 60 % higher than the next most impactful species, with increases of up to 520 % for scrub and 382 % for woodland (Fig. 7, Appendix C). Over the 100-year period, these higher interaction strengths led to significant reductions in woodland regeneration, pushing it toward the targeted values of 50 %, 45 %, and 25 %, versus the current projections of 60 % without the new species.

4. Discussion

In this paper, we applied a new EEM model to the Knepp Estate rewilding project, a case study characterised by a relatively high level of data availability, drawing from both time-series abundance data and 'reality checks' derived from expert knowledge and ecological literature. Calibration is a documented challenge in EEMs (Peterson and Bode, 2021; Vollert et al., 2024b) and becomes increasingly difficult as the number of constraints increases. To address this, we use a new GA-ABC approach. Our results show that while this method can efficiently search a large parameter space and optimise multiple constraints from diverse data sources, further improvements in EEM calibration techniques are necessary, especially for systems like Knepp which are too data-limited for system identification based purely on optimisation outputs, but too data-rich for simple random search methods.

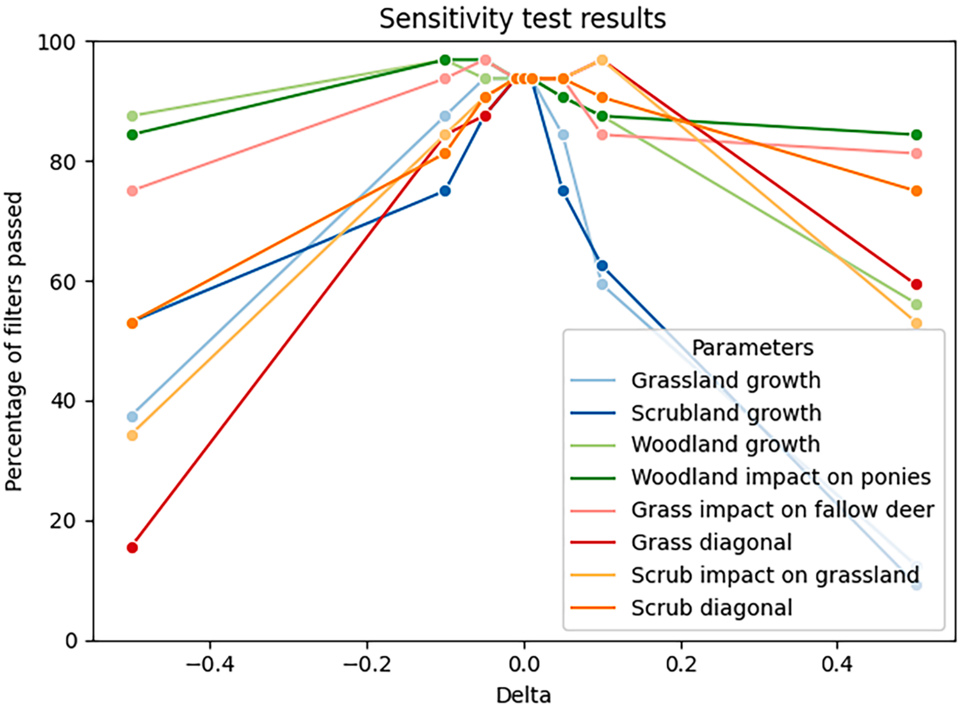

Fig. 4. Outputs of a one-at-a-time sensitivity analysis, in which all parameters are varied ±50, ±10, ±5, or ±1 % (delta) around their values, and the percentage of constraints passed recorded. This figure shows the top eight model parameters which had the highest impact on the percentage of passed constraints. A delta value of 0 shows the original parameter set obtained during the GA process, which passed 94 % of constraints.

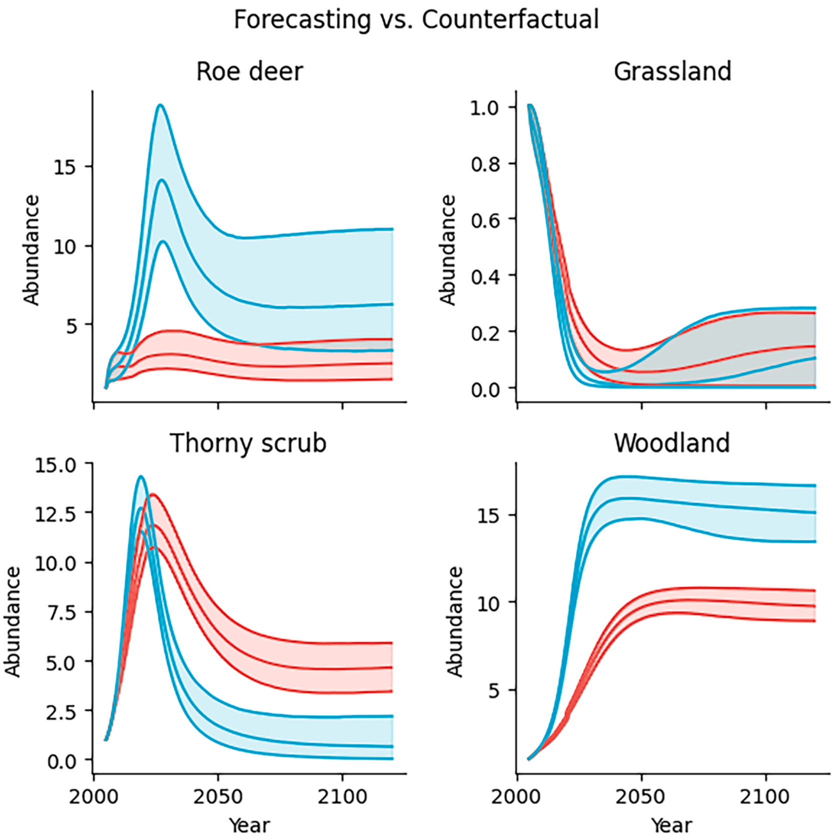

Fig. 5. A time-series of the Knepp ecosystem showing the accepted pre-reintroduction, post-reintroduction, and forecasting simulations. Red areas, which represent a total of 2045 runs, show the outcomes of current dynamics, with the assumption that all reintroduced species will be culled back to their 2020 stocking densities at each time point in the forecasting simulation. Blue areas also represent 2045 runs, and show what may have occurred if reintroductions had not taken place. Each graph shows the median and 95th percentile (shaded areas) for the four nodes which were naturally present pre-reintroductions and have no ongoing management via culling. All initial conditions are rescaled to 1 in the model. For the unmanaged species, initial conditions were 89.9 % grassland, 4.3 % scrubland, 5.8 % woodland, and 12 roe deer. An output value of 10 woodland, for example, is therefore equivalent to approximately 58 %.

4.1. An assessment of the GA-ABC method

The new GA-ABC method offered several advantages. The GA efficiently explored the large parameter space and optimised multiple nonlinear constraints simultaneously, identifying outputs with good fits to the time-series data. By manually adjusting the parameters to narrow the search space, we were able to identify a set that met both the abundance time-series constraints and all ’reality checks.’ This suggests the GA approach is a promising option for navigating the large parameter spaces typical of EEMs, especially when multiple constraints are used, and random searches are insufficient. The coupling of GAs with ABC added further value by identifying additional suitable parameters and offering probabilistic insights into uncertainty, improving the robustness of model predictions and providing a range of outcomes.

Despite these advantages, our approach faced some limitations. The need for manual trial-and-error adjustments to the GA parameter bounds was time-intensive and reduced the intended efficiency of the GA compared to other EEM calibration approaches (e.g., Baker et al., 2019). Expert knowledge and insights from ecological literature can help with defining realistic parameter bounds and potentially reduce the need for exploratory trial-and-error adjustments, but this is not always feasible in ecosystem models like ours, where many parameters are unknown. Passing all habitat-related constraints was particularly challenging, requiring extensive manual adjustments that often resulted in narrow GA bounds. This difficulty may have stemmed from the complex interactions between habitat types, such as light competition and woodland facilitation by thorny scrub, which are likely harder for the GA to resolve than simpler trophic interactions. Additionally, our objective function targeted the centres of constraint bounds, potentially causing the GA to overemphasise pushing certain nodes towards these exact centres even if they were already within acceptable bounds. This may have created trade-offs where certain species or dynamics were prioritised, making it harder to achieve a balanced solution that satisfied all constraints. A more flexible approach could involve adjusting the objective function to penalise only when nodes fall outside acceptable ranges, rather than for any deviation from the centre.

After manual adjustments, only 1 in 100 GA outputs passed all constraints. Ideally, our model would have been calibrated with many different outputs, leading to a diversity of solutions. To increase variability, we constructed a broad uniform prior distribution ranging ±10 % around the single acceptable output for the ABC process. This saved time and computational resources but came with several trade-offs. First, relying on a single GA output makes our outputs highly dependant on this specific solution, potentially complicating the reproducibility of our findings. Second, exploring a limited range around a single GA output may result in an incomplete understanding of the system’s possible variability and behaviour and could underestimate uncertainty in model outcomes. In our case, the ensemble of models consequently exhibited similar trajectories rather than a wide range of dynamics. Incorporating random noise into the process, as per Baker et al. (2019), could increase output variability. Running GAs many thousands of times could also provide ≈100 acceptable and diverse outputs, although this would be computationally expensive.

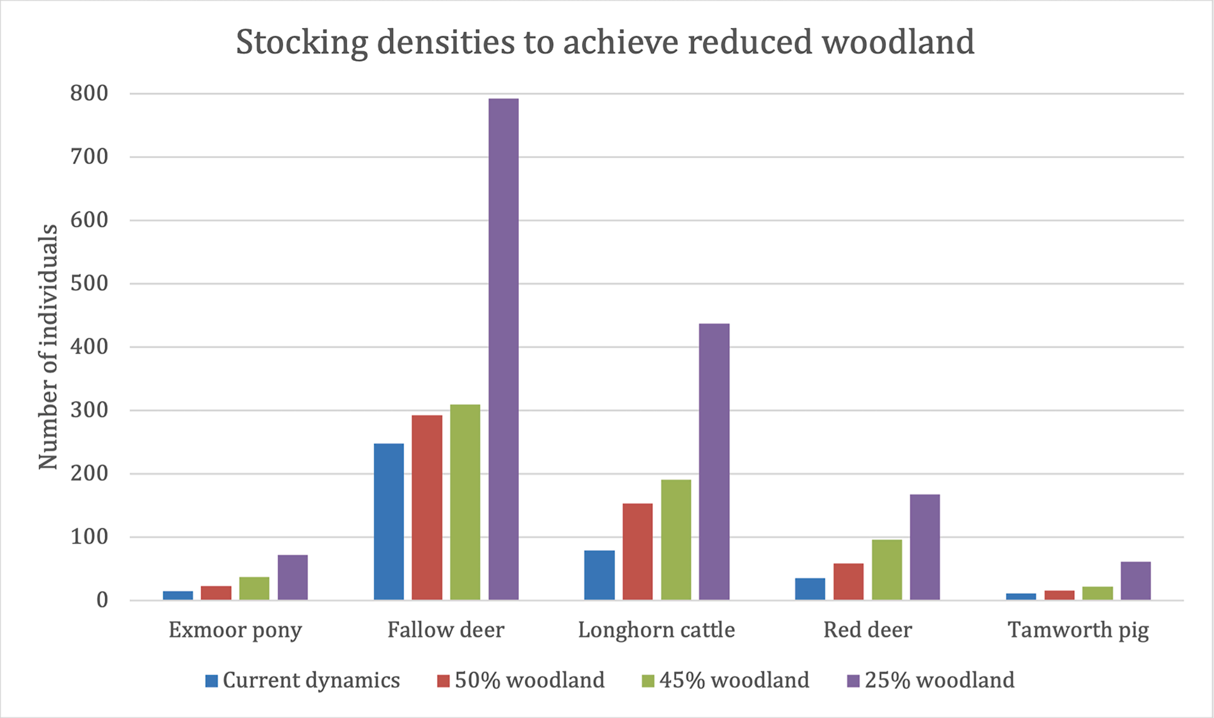

Fig. 6. A bar graph showing the collective change in stocking densities (number of individuals per species) required to achieve a given percentage reduction in woodland area, compared to the stocking densities used in 2020.

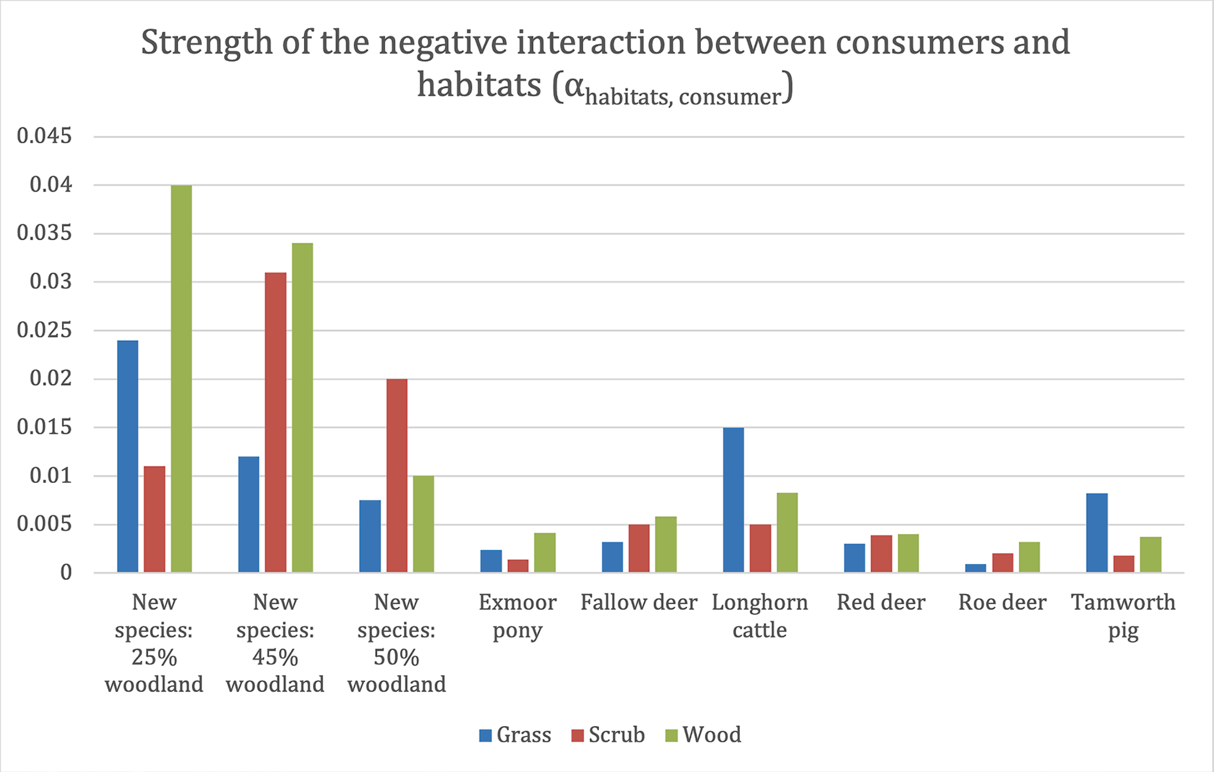

Fig. 7. A bar graph depicting the strength of the negative interactions between a hypothetical new species and each of the three habitat types $\alpha_{\text{habitats, consumer}}$ to achieve reduced woodland, compared to the impacts of other consumers in the absence of the new species. To achieve a reduced woodland area in future, the new species would generally need to have higher impacts on habitats than the other consumers.

4.2. Next steps for identifying additional parameter sets

There are several promising next steps to help identify a wider range of acceptable parameter sets. First, alternative ensemble generation methods could improve the efficiency of parameter space exploration. Particle swarm optimisation, for example, has shown improved convergence speed over GAs (Panda and Padhy, 2008), and sequential Monte Carlo sampling offers greater computational efficiency than existing EEM methods (Vollert et al., 2024b).

Incorporating additional complexities in interactions could also improve model fit. For instance, Type II functional responses - commonly observed in herbivores (Gross et al., 1993) – might better capture the consumer-resource dynamics in this system, although this would substantially increase the number of parameters. Another potential improvement involves integrating spatially dependant interactions, which could improve realism and better fit the available data. Specific habitat features like thorny scrub might reduce herbivory pressure on adjacent woodland, whereas grassland could have the opposite effect by increasing exposure of tree saplings to consumers (Van Uytvanck et al., 2008). Additionally, differences in consumer habitat preferences may lead to spatially distinct population concentrations, further influencing interaction strengths. By abstracting these spatial dynamics, the existing model might oversimplify how consumers and habitats interact, making it harder to replicate the real-world data. While incorporating these factors would come with the disadvantages of a more complex model, it may have more success passing the constraints.

Finally, a key decision we made was to set consumer self-interaction strengths $\alpha_{i,i}$ and negative growth rates (amounting to twelve parameters in this model) through fixed logic. This minimised the search space but may have inadvertently increased the difficulty of finding acceptable solutions. For example, by setting zeroes on the diagonal for consumers, we removed inherent population self-regulation, which could have destabilised the system and made it harder to pass certain constraints. Although explicit stability and feasibility requirements were not included in our model, maintaining system stability could have had advantages. Constraint 2, which required specific long-term trajectories for woodland and thorny scrubland over 100 years, was particularly difficult to satisfy, with no accepted runs prior to manual adjustments. The 100-year timeframe likely aligns with the system reaching a steady or pseudo-steady state, so by imposing stability criteria, we could have encouraged more balanced and consistent habitat dynamics and helped to satisfy this longer-term constraint. An alternative approach to fixing these parameters would be to set negative diagonals for consumers while setting $r_i = 0$. This would help maintain system stability while also ensuring consumers decline in the absence of food, potentially leading to the identification of more acceptable parameter sets.

4.3. Practical applications for ecological management

Overall, our results reinforce the findings of previous EEM studies that this methodology can help to guide conservation management and decision-making. Our EEM outputs support the idea that naturalistic grazing with large herbivores can drive the formation of habitats with multiple vegetation types. In all models, the reintroduction of large consumers is the driving factor in maintaining a grassland-scrubland-woodland ecosystem. Roe deer alone were insufficient to prevent significant woodland regeneration, with simulations indicating a woodland-dominant ecosystem within fifty years without further reintroductions.

The EEM also shows promise for guiding targeted management interventions. Our stocking density experiments, for example, optimised consumer populations to inform the levels needed to limit woodland expansion in the long-term, with results supporting studies that link increased grazing with reduced tree regeneration (Kinnaird, 1974; Pollock et al., 2005). Outputs from our second optimisation experiment suggest a hypothetical new species would generally have a greater negative impact on habitats than other consumers to reduce woodland abundance, demonstrating the model's utility in determining species characteristics needed to achieve specific goals. These optimiser outputs represent just one of potentially many combinations that could achieve the goal of reduced woodland. As such, they should be interpreted with caution and treated as guides for decision-making, rather than definitive targets. Variable bounds can be adjusted if management is interested in particular stocking ranges or herbivory impact levels.

5. Conclusion

Our study demonstrates that EEMs can successfully integrate a relatively high-level of data from multiple sources to replicate observed ecosystem dynamics and potentially inform management interventions. The GA-ABC approach shows promise for efficiently searching large parameter spaces to identify acceptable models, particularly in systems which have many constraints, but further improvements are needed. Using more efficient methods for ensemble generation, such as particle swarm optimisation or sequential Monte Carlo sampling approaches (Vollert et al., 2024b), or potentially developing a more complex model within the EEM framework, offer promising paths forward. Ultimately, refining the EEM methodology to improve calibration and facilitate the integration of more real-world data will enhance its utility in ecosystem management efforts, helping to achieve more effective and informed outcomes.

CRediT authorship contribution statement

Emily Neil: Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Conceptualization. **Ernesto Carrella:** Writing – review & editing, Supervision, Conceptualization. **Richard Bailey:** Writing – review & editing, Supervision, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank Penny Green and the Knepp Estate for providing data and expert knowledge, and the two anonymous reviewers for their valuable feedback.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.ecolmodel.2024.110948.

Data availability

Data will be made available on request.

References

Abdel-Basset, M., Abdel-Fatah, L., Sangaiah, A.K., 2018. Chapter 10 - Metaheuristic algorithms: a comprehensive review. In: Sangaiah, A.K., Sheng, M., Zhang, Z. (Eds.), Computational Intelligence for Multimedia Big Data on the Cloud with Engineering Applications. Academic Press, pp. 185–231. https://doi.org/10.1016/B978-0-12-813314-9.00010-4.

Anderson, P., Radford, E., 1994. Changes in vegetation following reduction in grazing pressure on the National Trust's Kinder Estate, Peak District, Derbyshire, England. Biol. Conserv. 69 (1), 55–63. https://doi.org/10.1016/0006-3207(94)90328-X.

Baker, C.M., Bode, M., Dexter, N., Lindenmayer, D.B., Foster, C., MacGregor, C., Plein, M., McDonald-Madden, E., 2019. A novel approach to assessing the ecosystem-wide impacts of reintroductions. Ecol. Appl. 29 (1), e01811. https://doi.org/10.1002/eap.1811.

Baker, C.M., Gordon, A., Bode, M., 2017. Ensemble ecosystem modeling for predicting ecosystem response to predator reintroduction: ensemble ecosystem modeling. Conserv. Biol. 31 (2), 376–384. https://doi.org/10.1111/cobi.12798.

Barkham, P., 2018. The magical wilderness farm: raising cows among the weeds at Knepp. The Guardian. https://www.theguardian.com/environment/2018/jun/15/the-magical-wilderness-farm-raising-cows-among-the-weeds-at-knepp.

Bellman, R.E., 1961. Adaptive Control Processes. Princeton University Press.

Bode, M., Baker, C.M., Benshemesh, J., Burnard, T., Rumpff, L., Hauser, C.E., Lahoz-Monfort, J.J., Wintle, B.A., 2017. Revealing beliefs: using ensemble ecosystem modelling to extrapolate expert beliefs to novel ecological scenarios. Methods Ecol. Evol. 8 (8), 1012–1021. https://doi.org/10.1111/2041-210X.12703.

Bremermann, H.J. (1958). *The evolution of intelligence. The Nervous System as a Model of Its Environment* (Technical Report 1). Department of Mathematics, University of Washington.

Broughton, R.K., Bullock, J.M., George, C., Hill, R.A., Hinsley, S.A., Maziarz, M., Melin, M., Mountford, J.O., Sparks, T.H., Pywell, R.F., 2021. Long-term woodland restoration on lowland farmland through passive rewilding. PLoS ONE 16 (6), e0252466. https://doi.org/10.1371/journal.pone.0252466.

Carver, S., Convery, I., Hawkins, S., Beyers, R., Eagle, A., Kun, Z., Van Maanen, E., Cao, Y., Fisher, M., Edwards, S.R., Nelson, C., Gann, G.D., Shurter, S., Aguilar, K., Andrade, A., Ripple, W.J., Davis, J., Sinclair, A., Bekoff, M., Soulé, M., 2021. Guiding principles for rewilding. Conserv. Biol. 35 (6), 1882–1893. https://doi.org/10.1111/cobi.13730.

Fraser, A., 1957. Simulation of genetic systems by automatic digital computers. II: effects of linkage on rates under selection. Austral. J. Biol. Sci. 10, 492–499.

Gallagher, K., Sambridge, M., 1994. Genetic algorithms: a powerful tool for large-scale nonlinear optimization problems. Comput. Geosci. 20, 1229–1236. https://doi.org/10.1016/0098-3004(94)90072-8.

Greenaway, T. (2006). *Knepp castle estate baseline ecological survey*. English Nature Research Reports.

Gross, J.E., Shipley, L.A., Hobbs, N.T., Spalinger, D.E., Wunder, B.A., 1993. Functional response of herbivores in food-concentrated patches: tests of a mechanistic model. Ecology 74 (3), 778–791. https://doi.org/10.2307/1940805.

Hamblin, S., 2013. On the practical usage of genetic algorithms in ecology and evolution. Methods Ecol. Evol. 4 (2), 184–194. https://doi.org/10.1111/2041-210X.12000.

Harding, P.T., Rose, F., 1986. Pasture-Woodlands in Lowland Britain: A Review of Their Importance for Wildlife Conservation. Natural Environment Research Council, Institute of Terrestrial Ecology.

Holland, J.H., 1975. Adaptation in Natural and Artificial Systems. University of Michigan Press.

Jepson, P., 2019. Recoverable earth: a twenty-first century environmental narrative. Ambio 48 (2), 123–130. https://doi.org/10.1007/s13280-018-1065-4.

Johnson, B.E., Cushman, J.H., 2007. Influence of a large herbivore reintroduction on plant invasions and community composition in a California grassland. Conserv. Biol. 21 (2), 515–526. https://doi.org/10.1111/j.1523-1739.2006.00610.x.

Kinnaird, J.W., 1974. Effect of site conditions on the regeneration of Birch (Betula Pendula Roth and B. Pubescens Ehrh.). J. Ecol. 62 (2), 467–472. https://doi.org/10.2307/2258992.

Kirby, K.J., 2004. A model of a natural wooded landscape in Britain as influenced by large herbivore activity. Forestry 77 (5), 405–420. https://doi.org/10.1093/forestry/77.5.405.

Kurwald, L. (2017). *Does rewilding restore soil biodiversity and function?* MSc Thesis, Cranfield University. https://knepp.co.uk/wp-content/uploads/2022/01/Does-Rewilding-restore-soil-biodiversity-and-function-Laura-Koester-genannt-Kurwald.pdf.

LeBauer, D.S., Wang, D., Richter, K.T., Davidson, C.C., Dietze, M.C., 2013. Facilitating feedbacks between field measurements and ecosystem models. Ecol. Monogr. 83 (2), 133–154. https://doi.org/10.1890/12-0137.1.

Lotka, A.J., 1925. Elements of Physical Biology. Williams and Wilkins.

Lovari, S., Ferretti, F., Corazza, M., Minder, I., Troiani, N., Ferrari, C., Saddi, A., 2014. Unexpected consequences of reintroductions: competition between reintroduced red deer and Apennine chamois. Anim. Conserv. 17 (4), 359–370. https://doi.org/10.1111/acv.12103.

Massart, D.L., Vandeginste, B.G.M., Buydens, L.M.C., De Jong, S., Lewi, P.J, Smeyers-Verbeke, J., 1998. Chapter 27 Genetic algorithms and other global search strategies. Data Handling in Science and Technology 20, 805–848. https://doi.org/10.1016/S0922-3487(97)80057-3.

Natural England. (2018). *Countryside stewardship: mid tier and new CS offers for wildlife manual*. https://assets.publishing.service.gov.uk/media/5c190598ed915d0c6bd78feb/mid-tier-and-offers-for-wildlife-manual-2018.pdf.

Navarro, L.M., Pereira, H.M., 2012. Rewilding abandoned landscapes in Europe. Ecosystems. 15 (6), 900–912. https://doi.org/10.1007/s10021-012-9558-7.

Neil, E., Carrella, E., & Bailey, R. (2024). *Knepp model.* github repository. https://github.com/emjneil/KneppModel.

Newton, A.C., Cantarello, E., Tejedor, N., Myers, G., 2013. Dynamics and conservation management of a wooded landscape under high herbivore pressure. Int. J. Biodivers. 2013, 1–15. https://doi.org/10.1155/2013/273948.

Nogués-Bravo, D., Simberloff, D., Rahbek, C., Sanders, N.J., 2016. Rewilding is the new Pandora's box in conservation. Curr. Biol. 26 (3), R87–R91. https://doi.org/10.1016/j.cub.2015.12.044.

Olff, H., Vera, F.W.M., Bokdam, J., Bakker, E.S., Gleichman, J.M., Maeyer, K.de, Smit, R, 1999. Shifting mosaics in grazed woodlands driven by the alternation of plant facilitation and competition. Plant Biol. 1 (2), 127–137. https://doi.org/10.1111/j.1438-8677.1999.tb00236.x.

Pace, M.L., Cole, J.J., Carpenter, S.R., Kitchell, J.F., 1999. Trophic cascades revealed in diverse ecosystems. Trends Ecol. Evol. (Amst.) 14 (12), 483–488. https://doi.org/10.1016/S0169-5347(99)01723-1.

Panda, S., Padhy, N.P., 2008. Comparison of particle swarm optimization and genetic algorithm for FACTS-based controller design. Appl. Soft. Comput. 8 (4), 1418–1427. https://doi.org/10.1016/j.asoc.2007.10.009.

Peterson, K., Bode, M., 2021. Using ensemble modeling to predict the impacts of assisted migration on recipient ecosystems. Conserv. Biol.: The J. of the Society for Conserv. Biol. 35 (2), 678–687. https://doi.org/10.1111/cobi.13571.

Pollock, M.L., Milner, J.M., Waterhouse, A., Holland, J.P., Legg, C.J., 2005. Impacts of livestock in regenerating upland birch woodlands in Scotland. Biol. Conserv. 123 (4), 443–452. https://doi.org/10.1016/j.biocon.2005.01.006.

Pritchard, J.K., Seielstad, M.T., Perez-Lezaun, A., Feldman, M.W., 1999. Population growth of human Y chromosomes: a study of Y chromosome microsatellites. Mol. Biol. Evol. 16 (12), 1791–1798. https://doi.org/10.1093/oxfordjournals.molbev.a026091.

Putman, R., 1986. Grazing in Temperate Ecosystems: Large Herbivores and the Ecology of the New Forest. Timer Press.

Raymond, B., McInnes, J., Dambacher, J.M., Way, S., Bergstrom, D.M., 2011. Qualitative modelling of invasive species eradication on subantarctic Macquarie Island. J. Appl. Ecol. 48 (1), 181–191. https://doi.org/10.1111/j.1365-2664.2010.01916.x.

Rendall, A.R., Sutherland, D.R., Baker, C.M., Raymond, B., Cooke, R., White, J.G., 2021. Managing ecosystems in a sea of uncertainty: invasive species management and assisted colonizations. Ecol. Appl. 31 (4), e02306. https://doi.org/10.1002/eap.2306.

Ross, P., Corne, D., 1994. Applications of genetic algorithms. AISB Q. on Evol. Comput. 23–30.

Ruxton, G.D., Beauchamp, G., 2008. The application of genetic algorithms in behavioural ecology, illustrated with a model of anti-predator vigilance. J. Theor. Biol. 250 (3), 435–448. https://doi.org/10.1016/j.jtbi.2007.10.022.

Scheffer, M., Carpenter, S., Foley, J.A., Folke, C., Walker, B., 2001. Catastrophic shifts in ecosystems. Nature 413 (6856), 6856. https://doi.org/10.1038/35098000. Article.

Schulte to Bühne, H., Ross, B., Sandom, C.J., & Pettorelli, N. (2022). Monitoring rewilding from space: the Knepp Estate as a case study. *J. Environ. Manag.*, 312, 114867. https://doi.org/10.1016/j.jenvman.2022.114867.

Smit, C., Ouden, J.D., Müller-Schärer, H., 2006. Unpalatable plants facilitate tree sapling survival in wooded pastures. J. Appl. Ecol. 43 (2), 305–312. https://doi.org/10.1111/j.1365-2664.2006.01147.x.

Smit, C., Ruifrok, J.L., 2011. From protégé to nurse plant: establishment of thorny shrubs in grazed temperate woodlands. J. Veg. Sci. 22 (3), 377–386. https://doi.org/10.1111/j.1654-1103.2011.01264.x.

Smit, C., Ruifrok, J.L., van Klink, R., Olff, H., 2015. Rewilding with large herbivores: the importance of grazing refuges for sapling establishment and wood-pasture formation. Biol. Conserv. 182, 134–142. https://doi.org/10.1016/j.biocon.2014.11.047.

Svenning, J.-C., 2020. Rewilding should be central to global restoration efforts. One Earth. 3 (6), 657–660. https://doi.org/10.1016/j.oneear.2020.11.014.

Svenning, J.-C., Pedersen, P.B.M., Donlan, C.J., Ejrnæs, R., Faurby, S., Galetti, M., Hansen, D.M., Sandel, B., Sandom, C.J., Terborgh, J.W., Vera, F.W.M, 2016. Science for a wilder Anthropocene: synthesis and future directions for trophic rewilding research. Proc. Natl. Acad. Sci. 113 (4), 898–906. https://doi.org/10.1073/pnas.1502556112.

ten Broeke, G., van Voorn, G., Ligtenberg, A., 2016. Which sensitivity analysis method should I use for my agent-based model? J. Artif. Soc. Soc. Simul. 19 (1), 5.

Tree, I., 2018. Wilding: the return of nature to a British farm. Picador.

Tuomi, M., Stark, S., Hoset, K.S., Väisänen, M., Oksanen, L., Murguzur, F.J.A., Tuomisto, H., Dahlgren, J., Bråthen, K.A., 2019. Herbivore effects on ecosystem process rates in a low-productive system. Ecosystems. 22 (4), 827–843. https://doi.org/10.1007/s10021-018-0307-4.

Van Uytvanck, J., Maes, D., Vandenhaute, D., Hoffmann, M., 2008. Restoration of woodpasture on former agricultural land: the importance of safe sites and time gaps before grazing for tree seedlings. Biol. Conserv. 141 (1), 78–88. https://doi.org/10.1016/j.biocon.2007.09.001.

Vera, F., 2000. Grazing Ecology and Forest History. CABI Publishing.

Vollert, S.A., Drovandi, C., & Adams, M.P. (2024a). *Ecosystem knowledge should replace coexistence and stability assumptions in ecological network modelling* (arXiv:2405.00333). http://arxiv.org/abs/2405.00333.

Vollert, S.A., Drovandi, C., Adams, M.P., 2024b. Unlocking ensemble ecosystem modelling for large and complex networks. PLoS Comput. Biol. 20 (3), e1011976. https://doi.org/10.1371/journal.pcbi.1011976.

Volterra, V., 1926. Variazioni e fluttuazioni del numero d'individui in specie d'animali conviventi. Mem. Acad. Lincei 2.