Environmental Modelling & Software 45 (2013) 74–91

Contents lists available at SciVerse ScienceDirect

Environmental Modelling & Software

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

Nonlinearities in biodiversity incentive schemes: A study using an integrated agent-based and metacommunity model☆

J. Gary Polhill*, Alessandro Gimona, Nicholas M. Gotts

The James Hutton Institute, Craigiebuckler, Aberdeen AB15 8QH, United Kingdom

ARTICLE INFO

Article history:
Received 2 April 2011
Received in revised form
29 October 2012
Accepted 6 November 2012
Available online 23 December 2012

Keywords:
Agents
Biodiversity
Land-use
Policy
Nonlinearity

ABSTRACT

We report results from over 20,000 runs of a coupled agent-based model of land use change and species metacommunity model. We explored the effect of increasing government incentive to improve biodiversity, in the context of other influences on land manager decision-making: aspirations, input costs, and price variability. The experiments test the four kinds of policy varying along two dimensions: activity-versus-outcome-based incentive, and individual-versus-collective incentive. The results from the experiments using boundedly rational agents, and comparison with profit-maximisation reveal thresholds in incentive schemes, where a sharp increase in environmental benefit occurs for a small increase in incentive. Further, the context affects the level of incentive at which turning points occur, and the degree of effect. Variability in outcome can also change with incentive and context, and some evidence suggests that environmental benefits are not always monotone increasing functions of incentives. Intuitively, if the incentive signal is large enough, land managers will farm the subsidy; and if the subsidy does not exactly match desired landscape outcomes, deterioration in environmental benefits may occur for higher incentives. Our results, whilst they suggest that outcome-based incentives may be more robust than activity-based, also highlight the importance of context in determining the success of agri-environmental incentive schemes. As such, they lend theoretical support to schemes, such as the Scottish Rural Development Programme, that include a localised component.

"the matter promises to be even more complex and mysterious than was originally supposed"

(Sir Arthur Conan-Doyle, *The Sign of the Four*)

1. Introduction

The Convention on Biological Diversity (Article 11) requires the subscribing countries to "adopt economically and socially sound measures that act as incentives for the conservation and sustainable use of components of biological diversity". Intensive agriculture is a major source of biodiversity loss, due to habitat destruction and loss of heterogeneity (Meehan et al., 2010; Kwaiser and Hendrix, 2008; Benton et al., 2003; Hald, 1999), and therefore an excellent candidate for policy intervention. Besides its important intrinsic value, biodiversity matters in agro-ecosystems because it can influence the long term sustainable productivity (Carvalheiro et al., 2011; Tilman et al., 2002; Naeem et al., 1995; Tilman and Downing, 1994), and hence system resilience, especially in view of increasingly costly inputs ultimately derived from oil. The Scottish Government adopted the biodiversity 2010 agenda of reducing biodiversity loss from agriculture.

In designing interventions, many governments favour a mix of market-based approaches and regulatory policy measures. Market-based approaches, according to economic theory, are more cost-effective, allow a flexible response to price signals, and avoid biodiversity being seen as a liability rather than an opportunity. Under this ideology, conservation incentives to land managers may be offered as voluntary measures aimed at correcting market failures causing the loss of species and ecosystem services. The removal of perverse (from the perspective of biodiversity conservation) incentives leading to over-intensification, has also been advocated as a policy measure (see, for example, Polasky et al., 1997). However, price volatility has the potential to compound the impact of intensification on biodiversity. Such volatility is a feature of liberalised agricultural markets with important effects on farmers' income, and as a consequence, on biodiversity levels (as we will show). Careful design of incentives based on an understanding of the underlying biological system is therefore crucial for policy success in agricultural socio-ecosystems.

☆ Thematic Issue on Spatial Agent-Based Models for Socio-Ecological Systems.

* Corresponding author. Tel.: +44 844 928 5428; fax: +44 844 928 5429.
E-mail address: gary.polhill@hutton.ac.uk (J.G. Polhill).

1364-8152/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.envsoft.2012.11.011

Cost-effectiveness is an important metric of policy evaluation. This can be interpreted in two ways with different consequences for evaluation. Firstly, a set of policy measures can be considered more cost-effective than others, if the total cost needed to achieve a policy goal is lower than for other measures. This interpretation is useful when there is a specific conservation objective, such as the survival of an endangered species. The second interpretation is concerned with overall conservation output. In this case a set of policy measures is more cost-effective than others if it results in higher overall biodiversity for a given cost. This point of view is useful when policy makers want to maximise the conservation output for a given available budget. In this study our focus is on overall biodiversity, measured by species richness at the landscape scale.

We have developed FEARLUS–SPOMM (Gimona et al., 2011; Gimona and Polhill, 2011) to explore biodiversity incentive schemes, by coupling FEARLUS (Polhill et al., 2001; Gotts et al., 2003; more recent versions described in Polhill et al., 2008; Gotts and Polhill, 2009; Polhill et al., 2010b), an agent-based model of agricultural land use change, with an enhanced version of Moilanen's (1999, 2004) Stochastic Patch Occupancy Model. Moilanen's (1999, 2004) model is a metapopulation model, simulating the occupancy of a single species in a space of connected patches (Levins, 1968), We have enhanced this model to simulate multiple species and interactions among them, making SPOMM a metacommunity model (Gilpin and Hanski, 1991; the extra '*M*' in SPO*M*M stands for 'Metacommunity'). In previous work (Gimona and Polhill, 2011), we explored the robustness of biodiversity policy and agri-environmental incentives across several scenarios of land manager, government policy and environmental attributes using a small sample of values for the incentive amount. These results suggested that it might be interesting to explore the effect of more gradually increasing the incentive, with a view to examining the relationship between incentive amount and species richness in more detail. We therefore chose a subset of the scenarios in Gimona and Polhill (2011) with which to increase the sample size of the incentive amount in this paper. We show here some selected results from this exploration, which amounted to over 20,000 runs of the coupled models, with a view to revealing more about the potential relationships between incentives and species richness, highlighting some of the sensitivities of biodiversity to farmer agent attributes, incentive scheme design, and other drivers of farmer behaviour.

In agro-ecosystems biodiversity is influenced by local management and by the landscape structure, which is a product of the decisions of individual land managers. Agent-based modelling is a natural tool to model the human portion of such systems, and is particularly well suited to studying coupled human–natural systems (Hare and Deadman, 2004; Boulanger and Bréchet, 2005), because it allows an intuitive representation of the environment and the embedding of agents within it. However, such couplings are not necessarily straightforward. Matthews et al. (2005) summarise various challenges in coupling social and environmental models, noting the stress many authors, reflecting on experiences in the area, place on a consistent integrated ontology in the coupled whole. In FEARLUS–SPOMM, some of these issues have been addressed because FEARLUS and SPOMM operate at compatible levels of abstraction and spatio-temporal scales. SPOMM was also specifically designed to be coupled with FEARLUS, and the process involved the developer of the latter.

The style of modelling in FEARLUS has been described by Boero and Squazzoni (2005) as a 'typification': the model constitutes a theoretical construct "intended to investigate some properties that apply to a wide range of empirical phenomena that share some common features" (para 3.8). As such, it is contrasted with 'theoretical abstractions' (Boero and Squazzoni cite work on the prisoner's dilemma as an example (Axelrod, 1997; Axelrod et al., 2002)) and 'case-based models', designed to be fitted to a particular time and place and provide an explanation of some of the phenomena observed there. Dean et al.'s (2000) work on the Anasazi is one of the examples given. As another example of a typification, Boero and Squazzoni (2005) cite their own work on industrial districts (Squazzoni and Boero, 2002; Boero et al., 2004). Typifications tend to make use of qualitative information, theory and second-order data (e.g. stylised facts) in exploring a class of phenomena, as opposed to case-based models, which will require significant amounts of quantitative data to fit to a specific instance of such a class. The version of FEARLUS used here is based on qualitative research with farmers and key informants in northeast Scotland, key assumptions in the model being checked with interviewees (Polhill et al., 2010b). Similar arguments would apply to the SPOMM, which may also be deemed a typification in the domain of ecology, and hence to FEARLUS–SPOMM. Typifications are useful for exploring questions 'in principle' about the relationships among phenomena in a class of systems.

One of the earliest agent-based models of a coupled human–natural system is Lansing and Kremer's (1994) work on Balinese water temples. This model was validated on empirical data, and was successfully used to persuade policy-makers of the merits of the water temple system for managing pests and irrigation. More recent work includes Guzy et al. (2008), who use a spatially-explicit agent-based model to assess the impact of urban expansion into farmlands and forests under various land use policy scenarios, and Brady et al.'s (2012) empirical model of the effects on ecosystem services of reforms of the European Union's Common Agricultural Policy on marginal agricultural regions in Sweden. Scenario analysis is a popular way to use agent-based models of coupled human–natural systems; Lempert (2002) recommend the use of ensembles of scenarios to model possible futures to explore robustness, resilience and stability of alternative policies. Participatory modelling techniques are often used in the study of such systems to capture local knowledge and engage with key stakeholders and decision-makers (Voinov and Bousquet, 2010). Recent examples include Anselme et al.'s (2010) work on shrub encroachment impacts on biodiversity conservation in the French Alps, and Lagabrielle et al.'s (2010) participatory process to integrate ecological knowledge into spatial planning on Réunion Island. A scenario approach is used here because FEARLUS–SPOMM is a typification: comparing results from as wide a range of scenarios as feasible with the computational power available is one way to avoid over-reliance on a specific instance of the model that has not been fitted to a particular case study in the real world.

Parker et al. (2008) outline various ways in which human and environmental systems can be coupled within a model, which can be divided into open-loop and closed-loop categories. In open-loop categories, submodels are executed sequentially, with no feedback from one submodel to another. Closed-loop categories feature such feedbacks, and although more challenging to implement, are clearly better fitted to capturing the complexity of the co-evolving landscape. An et al.'s (2005) IMSHED model, for example, is able to explore responses of households and panda populations to different conservation scenarios; they argue that the inclusion of feedbacks in their model leads to more representative results. Manson's (2005) SYPRIA model also features closed-loop human–environment interactions. FEARLUS–SPOMM currently features closed-loop interactions when outcome- rather than activity-based biodiversity incentive schemes are used.

2. Method

Here we give an overview of FEARLUS–SPOMM using Grimm et al.'s (2006; 2010) ODD (Overview, Design concepts and Details) model description protocol, with a slight modification to introduce the scenarios used as part of the overview. In subsection 2.2, we outline the simulations conducted, and in 2.3 the method used to interpret the results. Note also that here we describe FEARLUS–SPOMM as configured for these experiments. Other configurations are possible.

2.1. A coupled human and natural system model

2.1.1. Overview

2.1.1.1. Purpose. The purpose of the model is to explore how effective *in principle* various agri-environmental incentive schemes are at managing catchment biodiversity in the context of a number of other influences on land manager decision-making.

2.1.1.2. Entities, state variables and scales. The key entities in the model are an environment, consisting of a toroidal grid of 25 × 25 land parcels (a.k.a. patches), each owned by a land manager, who chooses a land use (conceptually, a combination of crop and land management strategy) for each parcel they own every year (the time step in the model). Land managers use a satisficing algorithm to choose land uses; no change is made to the land uses unless their profit (return minus input costs) is below an aspiration threshold for a number of consecutive years specified by their 'change delay'. When changing land uses they use a simplified form of case-based reasoning (Aamodt and Plaza, 1994), searching their case base for prior experiences of a land use, consulting neighbours for advice or experimenting if no such experience is available. Each experience (or 'case') stores the land parcel and economy as the context in which the decision was made, the land use chosen as the action taken, and the profit made as the outcome.

Each land use has a different yield, and this together with an exogenous economy time series and break-even threshold parameter (representing input costs) determines the economic return from the market accumulated to the land manager's account. Each land use makes one or more habitats for species available where it is used. Each patch records the species living on it as presence or absence (i.e. numbers of individuals are not recorded). Each species has parameters affecting its dispersal distance (alpha) and patch extinction probability (mu). A government agent monitors the environment, and has a rule it uses to issue an incentive to land managers, designed to prevent biodiversity loss.

Land managers offer a fixed price for land parcels that come up for sale due to neighbours going bankrupt, once their account exceeds their 'land offer threshold'. Land managers are regarded as bankrupt when their account is negative. A UML class diagram of the salient features of FEARLUS–SPOMM is shown in Fig. 1.

2.1.1.3. Process overview and scheduling. Each time step (year), consists of:

(i) Land managers choose land uses for each of the land parcels they own.
(ii) The land uses are used to derive a habitat map for species.
(iii) Species compute patch extinction and colonisation probabilities based on the habitat map and their prior distribution.
(iv) The government agent uses its rule to issue financial incentives to land managers.
(v) An economic return from the market is computed from the yield of the land use, and the current state of the economy, and added to the land managers' accounts.
(vi) Land managers update their case base.
(vii) Land managers with a negative account are regarded as bankrupt, and sell their land parcels.

2.1.1.4. Scenarios. A scenario is defined by four components: the rule used for the government agent, the time series used for the market, the break-even threshold for land managers, and the aspiration level. For the sake of convenience, we will use a nomenclature to refer to scenarios, consisting of a string of the following format: '$g/m/b/a$', where g refers to the government rule, m to the market, b to the break-even threshold (one of {25, 30}), and a to the aspiration threshold (one of {1, 5}). All scenarios use the same setup for species, habitats and land uses, which are the same as those in Gimona and Polhill (2011).

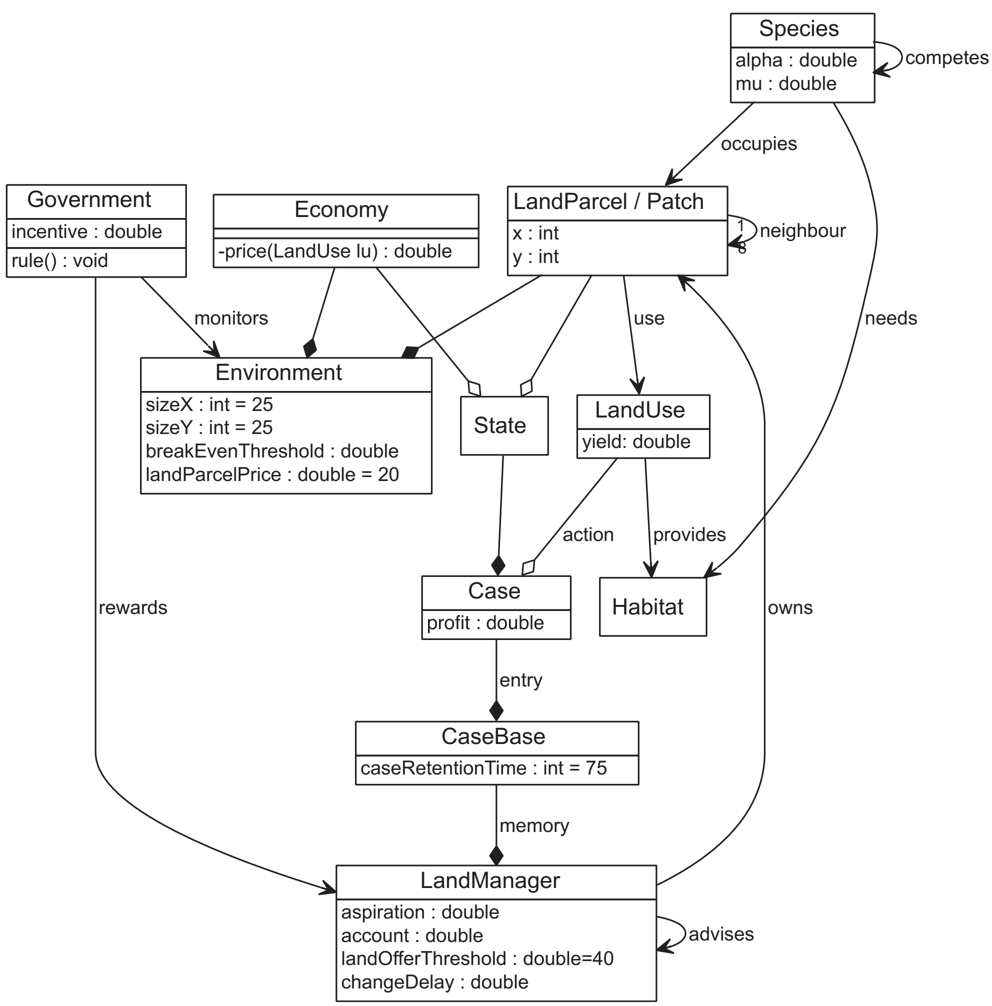

Fig. 1. UML class diagram showing the structure of FEARLUS–SPOMM.

The experiments in this paper use six land uses comprised of two classes having three levels of intensity: GL1, GL2 and GL3 (from low to high intensity; yielding 4, 5 and 6 in arbitrary units of yield respectively) representing grazing; AL1, AL2 and AL3 (low to high intensity; yields 4.5, 5.5, 6.5) representing arable. See Fig. 2. Land uses are mapped onto habitat provision: AL1 provides AH1, AL2 provides AH2, AL3 provides AH3, GL1 provides GH1, GL2 provides 20% GH1 and 80% GH2, GL3 provides GH3.

There are ten species, with parameters in Table 1; all can survive on at least one of habitats AH1 or GH1, and all except a competitor species (C), which causes the local extinction of G1, G2 and G3 after three consecutive years of occupancy, can survive on at least one of AH1 or GH2. (For this reason, GL2 provides 20% GH1.) Species A2, A3 and G4–6 cannot survive on the habitats associated with the most intensive land uses. Thus the lower intensity land uses AL1, GL1 and GL2 are the most important for biodiversity, but they have the lowest yields. Species G1–4 can only survive on grazing land uses, whilst A1 can only survive on arable. Species A1, C and G4–6 have the smallest dispersal kernel (which is inversely related to α – see Equation (1) below), whilst A3 and G1 have a larger dispersal kernel. Note also that GL1, the lowest intensity 'grazing' land use provides the best habitat for species C—a setup designed to implement one of many mechanisms, from real-world findings such as those in Wallis de Vries et al. (1998), by which biodiversity can be improved through grazing. This is consistent with the 'intermediate disturbance hypothesis' (Connell, 1978) in ecology: that biodiversity first rises, then falls, as the frequency and intensity of disturbance to an environment increases. To obtain a maximum landscape scale species richness of 10, a mixture of lower intensity land use types is

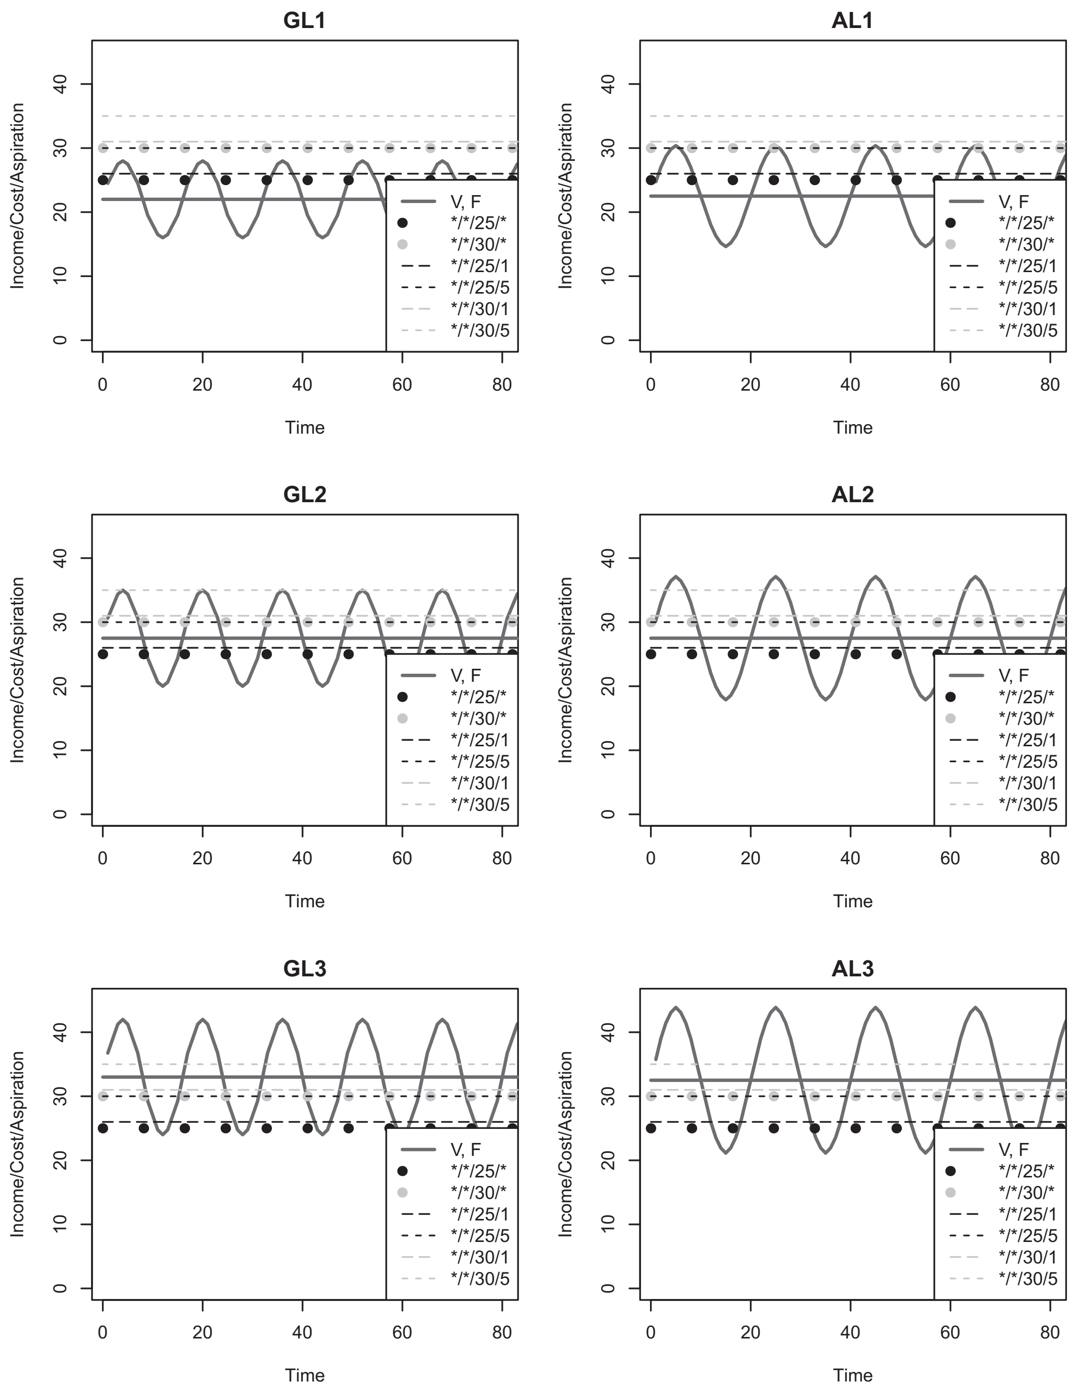

Fig. 2. Income from each land use in each economy (solid lines—the sinusoidal solid line is the 'variable' economy (V); the constant solid line is the 'fixed' economy (F)), compared with break even thresholds (dotted lines: dark dotted line 25, light 30) and break-even threshold plus aspiration thresholds (dashed lines: long dashes aspiration threshold 1, short dashes 5; dark when applied to break-even threshold 25, light for 30).

Table 1
Parameters for species used.

Spp.	α	μ	Habitats						Competition (time)
			GH1	GH2	GH3	AH1	AH2	AH3	
A1	1.3	0.1	–	–	–	Y	Y	Y	None
A2	0.9	0.1	–	Y	–	Y	Y	–	None
A3	0.8	0.1	–	Y	–	Y	–	–	None
C	1.3	0.05	Y	–	–	–	–	–	G1, G2, G3 (3)
G1	0.8	0.1	Y	Y	Y	–	–	–	None
G2	0.9	0.1	Y	Y	Y	–	–	–	None
G3	1.1	0.1	Y	Y	Y	–	–	–	None
G4	1.3	0.1	Y	Y	–	–	–	–	None
G5	1.3	0.1	Y	Y	–	Y	–	–	None
G6	1.3	0.1	Y	–	–	Y	–	–	None

required, including GH2 (or possibly GH3) to provide a refuge for G1–3 from C, without too much fragmentation so that species can re-disperse to patches with suitable habitat on which they become extinct.

We tried four rules for the incentive scheme, each intended to lessen the impact of the competitor species C (Gimona and Polhill, 2011, p. 179). Each rule uses an incentive computed from two parameters, reward (runs used integers in [1, 10]) and ratio ({1, 2, 10} tried), as follows (parentheses refer to the government rule setting in the scenario nomenclature described above):

- *Activity* rule (g = 'A'): Pay reward to each land manager for each parcel using GL2 or AL1. Note that rewarding for GL2 instead of GL1 is intended to avoid incentives that will cause the extinction of G1–3 through competition with C.
- *Outcome* rule (g = 'O'): Pay reward to each land manager for each occupancy of A2, A3, G3, G5 or G6 on a parcel they own. G3, with the smallest dispersal kernel, is the most vulnerable of G1–3. Rewarding for presence of G3 is intended to provide an incentive for land managers to provide a refuge for these three species.
- *Cluster-Activity* rule (g = 'CA'): Pay reward/ratio to each land manager for each parcel using GL2 or AL1, plus reward/ratio for each (Moore) neighbouring parcel using the same land use.
- *Cluster-Outcome* rule (g = 'CO'): Pay reward/ratio to each land manager for each occupancy of A2, A3, G3, G5 or G6 on a parcel of land they own, plus reward/ratio for each neighbouring parcel also having the same species.

We define 'incentive' as the amount paid by the rule. As shown above, for Activity and Outcome rules, the incentive is equal to the reward parameter; for Cluster-Activity and Cluster-Outcome rules, the incentive is equal to the reward parameter divided by the ratio parameter.

In our experiments, we used two time series for the economy, one having a fixed price for each land use, the other a variable price. The 'fixed' market (m = 'F' in the scenario nomenclature) offered 5 units of wealth per unit yield of AL1, AL2 and AL3, and 5.5 per unit yield of GL1, GL2 and GL3. The 'variable' market (m = 'V') oscillated approximately sinusoidally between 3.25 and 6.75 with period 20 years for AL1, AL2 and AL3; and between 4 and 7 with period 16 years for GL1, GL2 and GL3. A regular oscillation was used rather than (for example) a random walk, to make the effects of variability easier to discern: note that the land manager agents lack the intelligence to learn temporal patterns.

Fig. 2 compares the income from each land use with break-even thresholds and aspiration thresholds. In the 'fixed' market, GL1 and AL1 always make a loss; GL2 and AL2 make a loss if the break-even threshold is 30; and GL3 and AL3 always make a profit. However, land managers will not satisfice on GL2 and AL2 when the break-even threshold is 25 unless the aspiration threshold is 1, neither will they do so on GL3 and AL3 when the break-even threshold is 30 unless the aspiration threshold is 1. In the 'variable' market, each land use is sometimes profitable, and sometimes makes a loss, with the exception of GL1 when the break-even threshold is 30, which always makes a loss. (AL1 makes a very small profit in this case, but not enough to meet aspirations.) Aspirations, however, are measured at the farm scale as a mean over land parcels. Thus land uses that are not profitable and/or do not meet aspirations can be retained if the land manager has more than one land parcel, with the other parcels using land uses that generate sufficient income to compensate for the loss.

2.1.2. Design concepts

2.1.2.1. Basic principles. Case-based reasoning (Aamodt and Plaza, 1994) is an artificial intelligence technique based on evidence about expert practice in a range of occupations. An expert practitioner in any area where complex, context-dependent decisions are required (such as medicine, law, farming or design) will frequently approach a new problem by recalling examples of similar problems encountered in the past, the solutions attempted, and how successful they were. Case-based reasoning requires the decision-maker to have an episodic memory – a record of specific past experiences against which to compare the current case. It involves selection of the most appropriate prior cases, comparison of the success encountered in prior cases treated in different ways, and adaptation of the details of any past solution suggested by this process to the current case. The case-based reasoning algorithm used here is simplified, lacking a rich matching algorithm that would be expected in an artificial intelligence implementation to determine the degree of analogy between the current situation and an experience in the episodic memory. Case-based reasoning nevertheless describes the principles on which parts of the decision-making algorithm are based. Other aspects of land manager decision-making draw on Simon's (1955) 'satisficing', and qualitative research of farmers in northeast Scotland (Polhill et al., 2010b), which, among other things, introduced the 'change delay' parameter for land managers.

Stochastic patch occupancy models such as Moilanen's (2004) SPOMSIM model the presence or absence of species on patches of land. We have extended the concept to model multiple species, addressing the concept of 'metacommunities' (groups of communities connected by dispersal) in ecology (Holyoak et al., 2005) through treating each cell in the modelled space as a landscape patch.

2.1.2.2. Emergence. The emergent outcomes from the model are the species richness, spatial distributions of land uses, and the ability of the land managers to stay in business.

2.1.2.3. Adaptation. Adaptation occurs through bankruptcy and in-migration of land managers with different settings for their change delay parameter. Where the change delay makes a difference to the ability of land managers to stay in business, those with less favourable settings will go bankrupt, with the effect that the distribution of change delay in the population of agents will be different from the distribution from which this parameter is initialised.

2.1.2.4. Objectives. Objectives are implicit, but land managers' decision-making algorithms are aimed at keeping them in business, and the government agent's object is to improve biodiversity.

2.1.2.5. Learning. Land manager agents learn by storing new cases (each experience they have of using any particular land use) in their case base.

2.1.2.6. Prediction. Land managers predict the economic return they expect to get from a particular land use choice on the basis of the cases stored in their case base. Where they have no case for a particular land use, they may base their prediction on the experience of a neighbouring land manager who is willing to give them 'advice' in the form of a case they have encountered and stored. Where they have no access to a case for a land use, either in their case base, or in the form of advice, they 'predict' that the untried land use will meet their aspiration threshold.

2.1.2.7. Sensing. Land managers are aware of the state of the economy, their experiences of using different land uses, and of their neighbouring land managers. The government agent knows whatever it needs to in order to implement the incentive rule. In the case of 'Activity', the government agent is aware of all land uses applied, in the case of 'Outcome', it is aware of the presence or absence of each species on each patch of land.

2.1.2.8. Interaction. Land managers interact with each other through offering and requesting 'advice' – experience with using land uses they have not used, or do not remember having used themselves. The government agent interacts with land managers by issuing the incentive in accordance with its rules. Land managers interact with species by choosing land uses making differing amounts of habitat available for them. The government agent, for the 'Outcome' and 'Cluster-Outcome' rules, interacts with species through observing their presence or absence.

2.1.2.9. Stochasticity. For species, stochasticity is used to determine patch level absence or presence, based on the formulae calculating probabilities of local extinction or occupation, as appropriate. In land managers, stochasticity is used to decide between two land uses with the same expected outcome.

2.1.2.10. Observation. Each time step, land uses, bankruptcies, land manager income, and government expenditure were recorded. Species occupation data were recorded every 10 time steps.

2.1.3. Details

2.1.3.1. Initialisation. The landscape is initialized to a random distribution of the land uses AL1 and GL1 chosen with equal probability on each parcel. All species are then allocated to those patches on which they can survive. Land managers are randomly assigned one parcel each, and have 0 initial account; their change delay is taken from a uniform integer distribution in the range [1, 9]. Aspirations are set according to the scenario the run belongs to (i.e. 1 or 5), as is the break-even threshold (25 or 30).

2.1.3.2. Input. The economy time series (one of 'fixed' (F) or 'variable' (V)) is the only input to the model.

2.1.3.3. Submodels

(i) Land use decision-making.

Each land manager computes the mean profit from all their land parcels in the previous year, and if this has been less than their aspiration for at least as long as their change delay (the number of consecutive years they are prepared to tolerate below-aspiration profits), they will use case-based reasoning to choose a land use for each land parcel they own; otherwise they will make no changes. Note that the enterprise-level aspiration means that there is no requirement for individual land parcels in a multi-parcel estate to fulfil aspirations.

Specifically, if A is the aspiration threshold of a land manager, M is the total economic return they received from the market in the previous year, P is the number of parcels they own, and R the amount of incentive the manager received in step (iv) of the previous year, then the aspiration threshold test above is that $(M+R)/P > A$.

To make a (potential) change using case-based reasoning, managers consider each parcel in turn, and for each land use calculate expected profit by searching the case base for an experience of using that land use in the expected state of the economy (a symbol standing for the relationship between land use and price at a particular time step), which is assumed to be the same as the previous year. If a case is found, the expected profit is that recorded in the case. If not, then if a neighbour[1] has a case for the expected state of the economy, the profit recorded there is used as the expected profit for the land use, otherwise the aspiration itself is used (i.e. it is assumed that the land use for which there is no experience will just meet aspirations). The land use assigned to the parcel is that with maximum expected profit; if there are two or more such land uses, one is selected at random.

The case base as used here may thus be conceived as a data structure consisting of a hash table for each land use, with key the symbol representing the state of the economy, and entry a (possibly empty) list of prices returned to the land manager when they have used the land use in that state of the economy. Searching the case base is a matter of looking up the list in the hash table for the land use in question, which is associated with the required economy state symbol. Where the list has more than one price, the most recent price is returned.

(ii) Update habitats.

The habitat map is computed from the land uses as described in *Initialisation* above.

(iii) Update occupancy.

Dispersal, colonization and local extinction are computed as per Moilanen (2004), with the extra property that if species C is present on a patch for four consecutive time steps, it causes local extinction of G1, G2 and G3 if present, and prevents them recolonising whilst occupancy continues. The Moilanen equations are given below:

Let $S_{ps}(t)$ be the connectivity of species s on patch p at time t, computed as:

$$S_{ps}(t) = A_{ps}^{c} \sum_{q \neq p} O_{ps}(t-1)\exp(-\alpha_s d_{pq}) A_{qs}^{b} \quad (1)$$

where A_{ps} is the amount of habitat made available on patch p for species s (see *Scenarios*), q iterates over all patches other than p, $O_{ps}(t)$ is an occupancy indicator variable (1 if patch p is occupied by species s at time t, 0 otherwise), α_s is a dispersal parameter for s (see Table 1), d_{pq} is the Euclidean distance between patches p and q (assuming a toroidal spatial topology), and b and c are parameters (both 1).

Then the probability that s colonises a patch p it currently does not occupy is given by (2), where y is a parameter (set to 1):

$$C_{ps}(t) = [S_{ps}(t)]^2 / \left([S_{ps}(t)]^2 + y^2 \right) \quad (2)$$

and the local extinction probability of s on a patch p it occupies is given by (3), where x is a parameter (set to 1), μ is the mortality parameter of s (Table 1), and A_{ps} is as per (1):

$$E_{ps} = \mu_s / A_{ps}^{x} \quad (3)$$

(iv) Implement incentives.

Managers receive funds from the government agent according to the rules described above.

(v) Economic returns.

Updating the state of the economy from the economy file provides a market price p_i for each land use i. A yield, y_i for each land use is set as described in the *Scenarios* section. If e is the break-even threshold, then the economic return for each land parcel is $p_i y_i - e$.

(vi) Case base update.

The case base of a land manager is updated on a parcel-by-parcel basis, assuming an equal distribution of reward across the parcels the land manager owns. A case is added for each parcel, containing the state of the economy, the land use chosen, and the economic return plus distributed reward: $p_i y_i - e + R/P$, where R is the total reward received by the manager in step (iv), and P the number of parcels they own. Any cases from more than 75 time steps ago are removed from the case base.

(vii) Land exchange.

If a land manager's account is less than zero, they are bankrupt, and all their land parcels are sold. For each parcel for sale, choose a new land manager at random from a set consisting of one new in-migrant land manager, and all land managers in the Moore neighbourhood of the parcel with an account >40. Assign the parcel to the manager selected. If the new owner is the in-migrant manager, then initialize its other variables (as per initialization, above); if not, deduct 20 from the account of the new owner. Note that in-migrant managers cannot be assigned more than one of the parcels for sale.

2.2. Run setup summary

Table 2 summarises the first set of parameter explorations done. Each combination of parameters (960 in total) was repeated 20 times using different seeds for the pseudo-random number generator, making 19,200 runs overall. (Note that, as discussed above, the ratio parameter has no effect on runs using government rules 'Activity' or 'Outcome'. Effectively, there were 60 rather than 20 replications of these parameter settings—this was done to provide sufficient samples for any statistical tests comparing the Cluster-*x* rules with their non-clustered counterparts.) Each run was for 200 time steps.

A second phase of parameter explorations was conducted for the Activity and Outcome government rules using higher values of Reward (see Table 3) to create runs with total expenditure more comparable to the levels seen in the Cluster-*x* rules. (Expenditure is somewhat complicated as the amount spent depends not only on the rule itself, but also on the success of the rule.) At 20 runs per parameter setting, this amounted to a further 2240 runs.

There are 32 scenarios in total, defined by the combinations of values explored for government rule, market, break-even threshold and aspiration threshold. Within each scenario, the reward and ratio parameters define the level of incentive applied by the government agent in accordance with its rule. In phase one, each scenario has 600 runs exploring different incentive settings; phase two adds 140 runs to the 16 scenarios A/*/*/* and O/*/*/* (using * as a wildcard) exploring further incentive settings.

Table 2
Summary of parameter exploration: first phase.

Parameter	Values
Government rule	Activity, Outcome, Cluster-Activity, Cluster-Outcome
Reward	1, 2, 3, 4, 5, 6, 7, 8, 9, 10
Ratio	1, 2, 10
Market	'Fixed', 'variable'
Break-even threshold	25, 30
Aspiration threshold	1, 5

[1] The neighbours of a land manager are those owning a parcel in the Moore neighbourhood of that land manager's parcels.

Table 3
Summary of parameter exploration: second phase.

Parameter	Values
Government rule	Activity, Outcome
Reward	15, 20, 25, 30, 40, 50, 100
Ratio	1
Market	‘Fixed’, ‘variable’
Break-even threshold	25, 30
Aspiration threshold	1, 5

2.3. Analysis

Runs were rejected from consideration if they led to a high bankruptcy rate (more than 10% of managers per year). Four hundred and eighteen runs fell into this category (all from scenarios matching */F/30/5, where there are no fluctuations in the market, and the demands on and ambitions of land managers are highest), leaving 21,022 runs. Runs were also rejected if they led to a very high expenditure by the government agent (mean annual expenditure over the last 100 time steps more than or equal to 25,000). This reduced the analysis by a further 4073 runs, mainly from runs using the Cluster-Outcome (2560) and Outcome (1120 – all the runs conducted in phase 2) government rules. This leaves 16,949 runs.

The remaining results were analysed with three points of interest in mind: (i) nonlinearity of the relationship between incentive and biodiversity outcomes; (ii) any effect of government rule on this relationship; (iii) effects of context (market, aspirations and break-even threshold). We fitted Generalised Additive Models (GAMs, Hastie and Tibshirani (1990)) to the data to examine trends, using species richness over the whole landscape as a measure of biodiversity. GAMs can be used to test for nonlinearity (e.g. Banks et al., 2000) by comparing fits generated using different numbers of smoothing terms. Nonlinearity can also be assessed by comparing a GAM fit with a fit using a linear model, using ANOVA to test if there is a significant difference between the fits, and comparing the sum of squared difference between the fitted functions to see if it is large. The Akaike Information Criterion (AIC; Akaike, 1973, 1974) is also used to compare statistical fits, and includes a penalty for the number of parameters. Lower values are preferred, and this can also be used to compare the linear model with the GAM.

To test (i), we deployed a number of tests of nonlinearity in the results data. We fitted a GAM *G*() to the data using thin-plate regression splines, and then compared it with two further statistical models: *K*(), a GAM where the dimension of the basis used to represent the smooth term was set to 4 (effectively constraining the fit to a cubic); and *L*(), a linear model. Using these three statistical models, the following tests were performed to assess the nonlinearity of each scenario:

1. *G*() has estimated degrees of freedom more than 3, and the *p*-value of the smooth term is significant.
2. An ANOVA comparing *G*() and *K*() has a significant *p*-value, the sum of squared difference between *G*() and *K*() is ‘large’, and the sum of squared error of *G*() on the data is less than that of *K*().
3. The AIC of *G*() is at least 2 less than that of *K*().
4. As per test 2, but comparing *G*() and *L*().
5. As per test 3, but comparing *G*() and *L*().

The GAM models were fitted to the relationship between *incentive* and landscape-scale species richness, where incentive is defined as reward/ratio for Cluster-*x* government rules, and reward otherwise. We used Wood’s (2001; 2006) MGCV package (version 1.7–2) for R to fit GAMs. This computes the smoothing parameters for *G*() by minimising an unbiased risk estimator criterion (see MGCV documentation). The sum of squared difference between two models, computed in tests 2 and 4, uniformly sampled each model at the same 10,000 points on the *x*-axis; a ‘large’ difference being inferred if the sum of squared difference is more than 10,000 – implying a mean difference in richness (the *y*-axis) of one species. Since 76 tests of significance are involved in the above tests on the 32 scenarios, the *p*-value for a significant result was set at 0.0001 to avoid type I error. The ANOVA and AIC were computed using standard R functions.

To examine (ii) and (iii) we used recursively partitioning classification trees (Breiman et al., 1984), with government expenditure, government rule, market, break-even threshold and aspiration as explanatory variables, and landscape-scale species richness as the response variable. We cannot use incentive as an explanatory variable for this analysis because of the differences in the ways each government rule uses this parameter. Expenditure, though it is itself an outcome variable, is a fairer comparator.

3. Results & analysis

Fig. 3 shows some of the results using a sunflower plot, in which the incentive and resulting species richness in each run is represented by a ‘petal’ of the sunflower (or a darker coloured small dot if there is only one run with that incentive/richness combination). Visual inspection of these results suggests there are thresholds in incentive occurring in scenario A/F/30/1 between 3 and 4 at which richness jumps from 4 to 8, and then from 8 to 9 between incentives 8 and 9. In scenario A/V/30/1, although the variable market has led to more uncertainty in the outcome, at incentives 4 and less species richness higher than 3 does not occur, whilst between 5 and 10 richness lies mainly in the range 4–7 (with a single run exception at incentive 5 having richness 3). A further threshold somewhere between incentives 10 and 15 results in modal

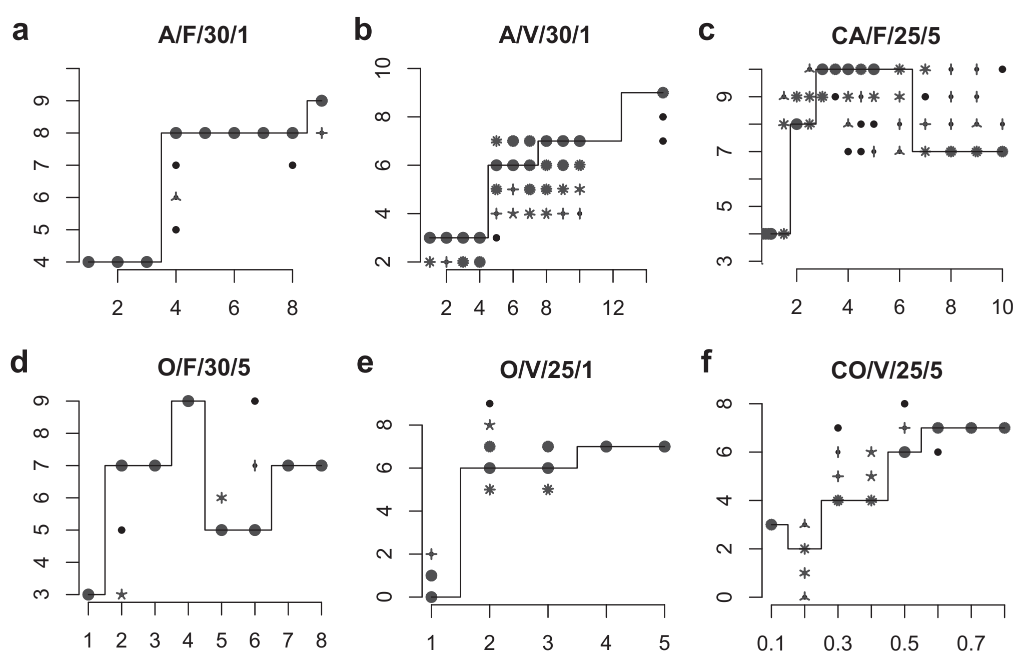

Fig. 3. Portions of the data from six of the scenarios using sunflower plots (see text), with incentive on the horizontal axis, and landscape scale species richness on the vertical. The prediction of a classification tree built using recursive partitioning on the incentive/richness relationship is shown using straight lines.

richness 9. Scenario CA/F/25/5 shows a threshold between incentives 1 and 1.5, where the richness jumps from 4 to 8 or 9; between incentives 6 and 10, higher richness values are seen less and less frequently as the modal richness declines from 10 to 7. Scenario O/F/30/5 has a peculiar shape – the richness rises from 3 to 7 somewhere between incentives 1 and 2, then to 9 between 3 and 4. This is followed by a decline to modal richness 5 between incentives 4 and 5, before a rise to 7 between incentives 6 and 7. Scenario O/V/25/1 has a clear threshold between incentives 1 and 2 where the modal richness jumps from 0 to 6. Scenario CO/V/25/5 has thresholds in the minimum richness as the incentive increases, from 3 to 0, then 4, 6 and 7 for incentives 0.1, 0.2, 0.3, 0.5 and 0.7 respectively.

The visual analysis is supported by predictions from a recursively partitioning classification tree fit to the incentive/richness relationship. Predictions from the tree are shown with straight lines in Fig. 3 – vertical lines show incentive partitioning thresholds, horizontal lines show the classification tree's predicted class, which where more than one richness value is associated with the same incentive, is based on the most common richness in runs over the incentive range. The classification tree identifies thresholds in incentive that are reasonably consistent with visual inspection of the raw data.

The results of the nonlinearity tests for each scenario are shown in Table 4, with graphs in Fig. 4 showing the fits provided by *G*(), *K*() and *L*() for scenarios selected for their relevance to analysis of Fig. 5. (The other scenarios are shown in Fig. A.1.) Twenty-eight scenarios passed at least two of the tests, and two scenarios passed all five. Scenario O/F/25/5 (Fig. 4(b), bottom left), one of the two scenarios passing none of the tests, did so because the GAM failed to fit the data, which, by inspection of the raw data plotted as sunflower plots, is nonlinear. In all other scenarios, the deviance in the raw data explained by *G*() was at least 68%; in 26 scenarios, the deviance explained was more than 80%. *G*() nevertheless tended to overfit the data (a common problem with GAMs), and in only six scenarios (one of which is shown in the top-right graph in Fig. 4(b)) does it provide a significantly better fit to the data than *K*() as measured by test 2. Even so, test 4 shows that *G*() is a significantly better fit than a linear model in all but nine scenarios (of which Fig. 4(c) shows an example). The AIC (test 5) seems to set a higher benchmark than test 4 for superiority of *G*() over *L*(), with ten scenarios passing this test. However, this is not so when comparing *G*() with *K*() (tests 2 and 3), where *G*() mostly has a better AIC than *K*(). Since the AIC includes a (negative log-) likelihood term as well as a penalty for the number of parameters, and *G*() typically has a higher number of parameters than *K*(), this suggests that, the results of test 2 notwithstanding, *G*() has a better enough likelihood than *K*() to overcome the parameter penalty difference.

The results in Fig. 5 show that expenditure is the most significant determinant of landscape-scale species richness: lower expenditure being generally associated with lower richness, and higher expenditure with higher richness. However, whether expenditure is high or low, all the scenario variables may have an effect on richness. Market variability has the most significant influence on richness when expenditure is low (split leading to leaf node A vs. B–H), and when expenditure is high, richness is most affected by the use of an activity- or outcome- (split to leaf node I vs. J–O) based government rule. This might be what one would expect: at some point as expenditure increases, government policy is a large enough proportion of land managers' income that it has a more significant effect on biodiversity than the market. Aspiration and break-even threshold influence richness in the fixed market when expenditure is low (nodes C–H); a clustering element to the activity-based government rule affects the distribution of richness in the variable market when expenditure is high (node K).

Nonlinearity in the expenditure/richness relationship is clear from this tree, particularly in the case of scenario A/F/25/5 (the incentive/richness relationship of which is shown in Fig. 4(a), in which the modal richness changes from 4, to 8, then 9, then 7 (nodes B and D, E and M, O, then N respectively) as the expenditure increases:

- Nodes B and D show a predominance of land use GL3, and some use of AL3: these will be the only land uses that satisfice for the incentive values associated with this low level of expenditure, and they only provide habitat for species G1–3 and A1. In a fixed market, GL2 returns a profit of 2.5 without any incentive. To meet the aspiration threshold in this scenario, an incentive of 2.5 is needed.
- Node E shows GL2 being used, with species G4–5 and A2–3, for which it provides habitat, showing low, but above zero levels of occupancy.
- Node M shows greater use of GL2 than node E, and use of AL1 in many of the runs, which requires an incentive of 7.5 to satisfice in this scenario. AL1 allows species G6 to survive, and increased use of GL2 also means greater abundance of G4–5 and A2–3.
- Node O has a landscape in which land managers are mostly using AL1 – despite GL2 having a higher profit – this is because

Table 4
Results of the nonlinearity tests for each scenario, with deviance explained by *G*() in parentheses. An asterisk is used to indicate that the conditions of the test were passed, a dash otherwise. See Section 2.3 for a description of the tests.

Scenario (Deviance explained)	Test 1	2	3	4	5	Scenario (Deviance explained)	Test 1	2	3	4	5
A/F/25/1 (68%)	*	–	–	–	–	O/F/25/1 (88%)	–	–	–	–	–
A/F/25/5 (83%)	*	–	*	*	–	O/F/25/5 (1%)	–	–	–	–	–
A/F/30/1 (85%)	*	–	*	*	*	O/F/30/1 (94%)	*	*	*	*	–
A/F/30/5 (88%)	*	*	*	*	–	O/F/30/5 (93%)	*	*	*	*	–
A/V/25/1 (77%)	*	*	*	*	*	O/V/25/1 (90%)	*	*	*	*	*
A/V/25/5 (95%)	*	–	*	*	–	O/V/25/5 (86%)	*	–	*	–	–
A/V/30/1 (83%)	*	–	*	*	*	O/V/30/1 (87%)	*	–	*	–	–
A/V/30/5 (87%)	*	–	*	*	–	O/V/30/5 (89%)	*	*	*	*	–
CA/F/25/1 (75%)	*	–	*	–	–	CO/F/25/1 (71%)	–	–	*	*	–
CA/F/25/5 (93%)	*	–	*	*	*	CO/F/25/5 (80%)	*	–	*	*	–
CA/F/30/1 (93%)	*	–	*	*	*	CO/F/30/1 (96%)	*	–	*	*	–
CA/F/30/5 (77%)	*	–	*	*	*	CO/F/30/5 (90%)	*	–	*	*	*
CA/V/25/1 (82%)	*	–	*	*	*	CO/V/25/1 (84%)	*	–	*	*	*
CA/V/25/5 (80%)	*	–	*	–	–	CO/V/25/5 (86%)	*	–	*	*	–
CA/V/30/1 (86%)	*	–	*	*	–	CO/V/30/1 (89%)	*	–	*	*	–
CA/V/30/5 (88%)	*	–	*	–	–	CO/V/30/5 (76%)	*	–	–	–	–

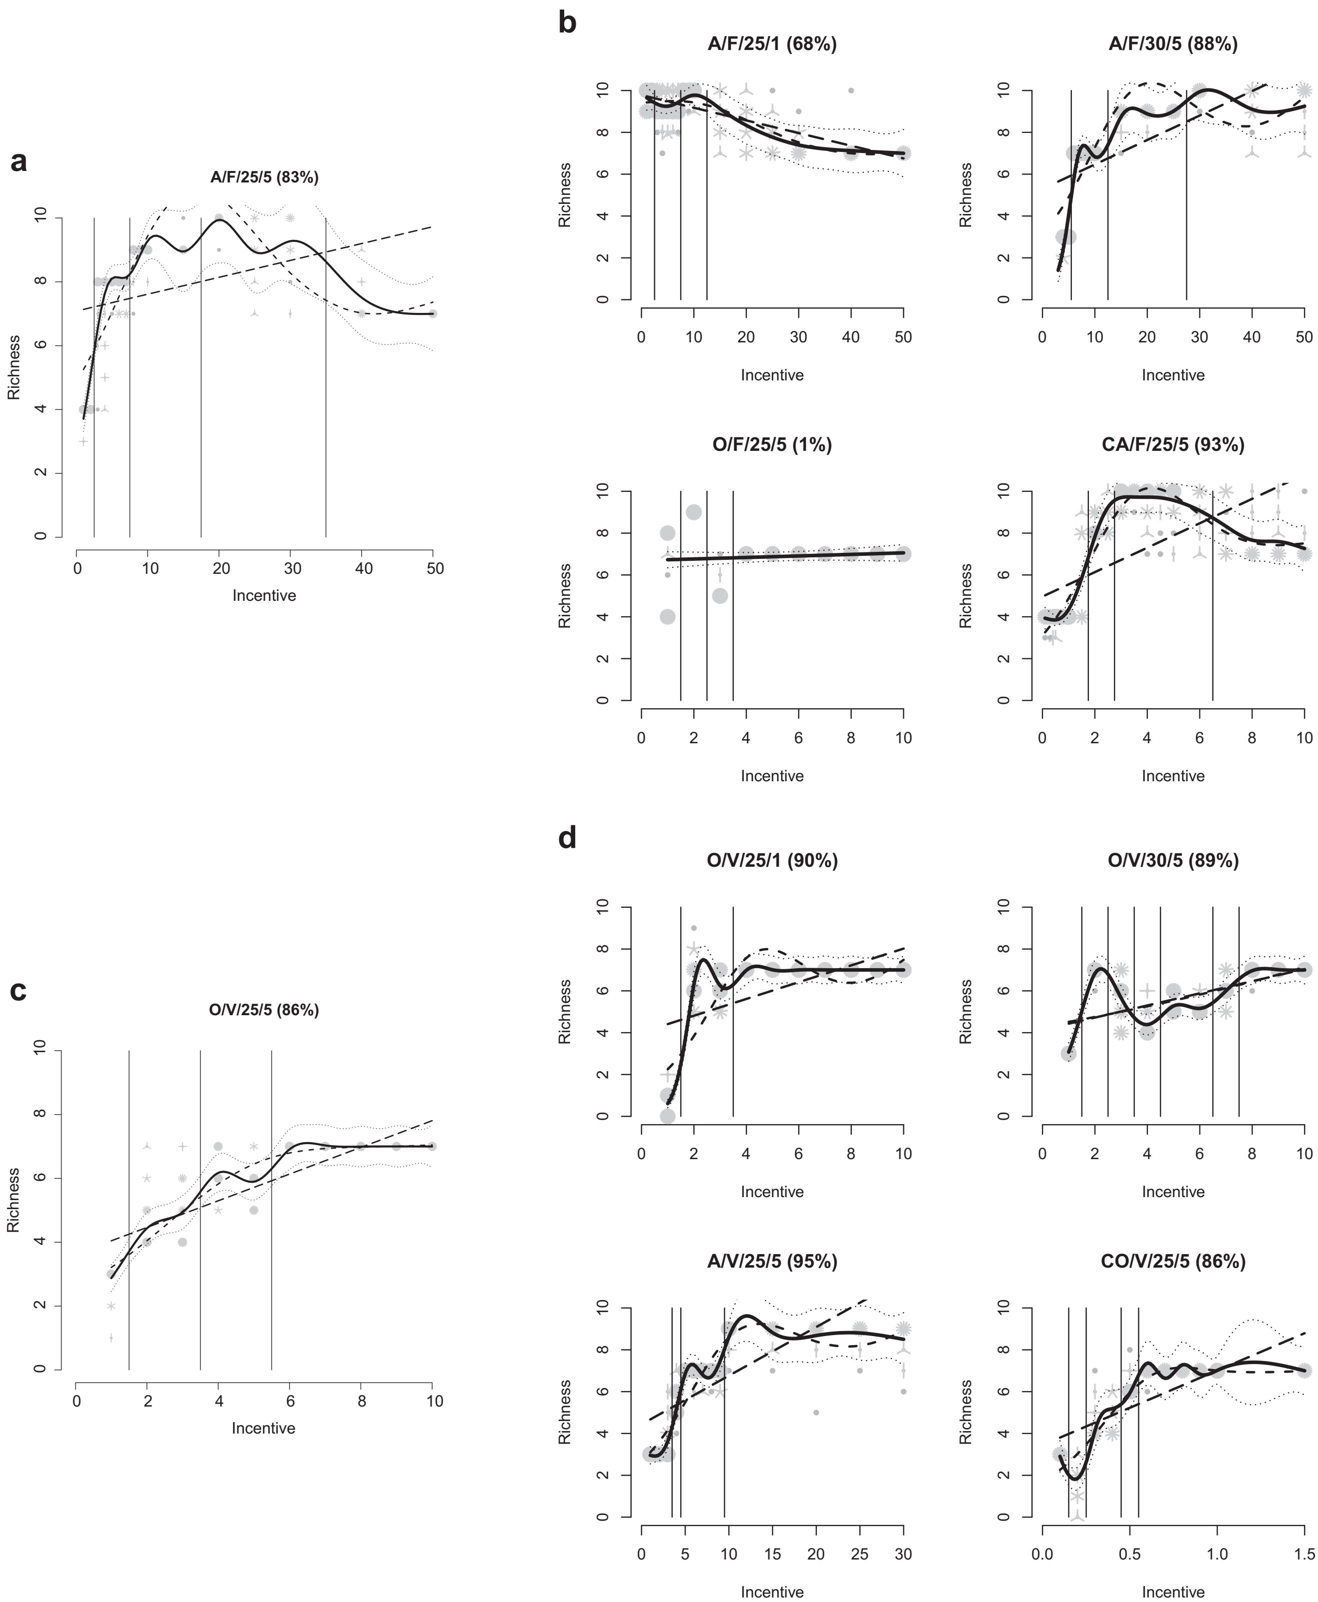

Fig. 4. Results from selected scenarios, shown as sunflower plots of incentive against richness. In each graph, *G*() is plotted with a thick black line with dotted lines showing 95% confidence intervals, *K*() with short dashes, and *L*() with long dashes. Vertical lines are drawn at split points identified by recursive partitioning of the incentive/richness relationship. Scenarios in (b) and (d) are related to those in (a) and (c) respectively by changing one scenario variable.

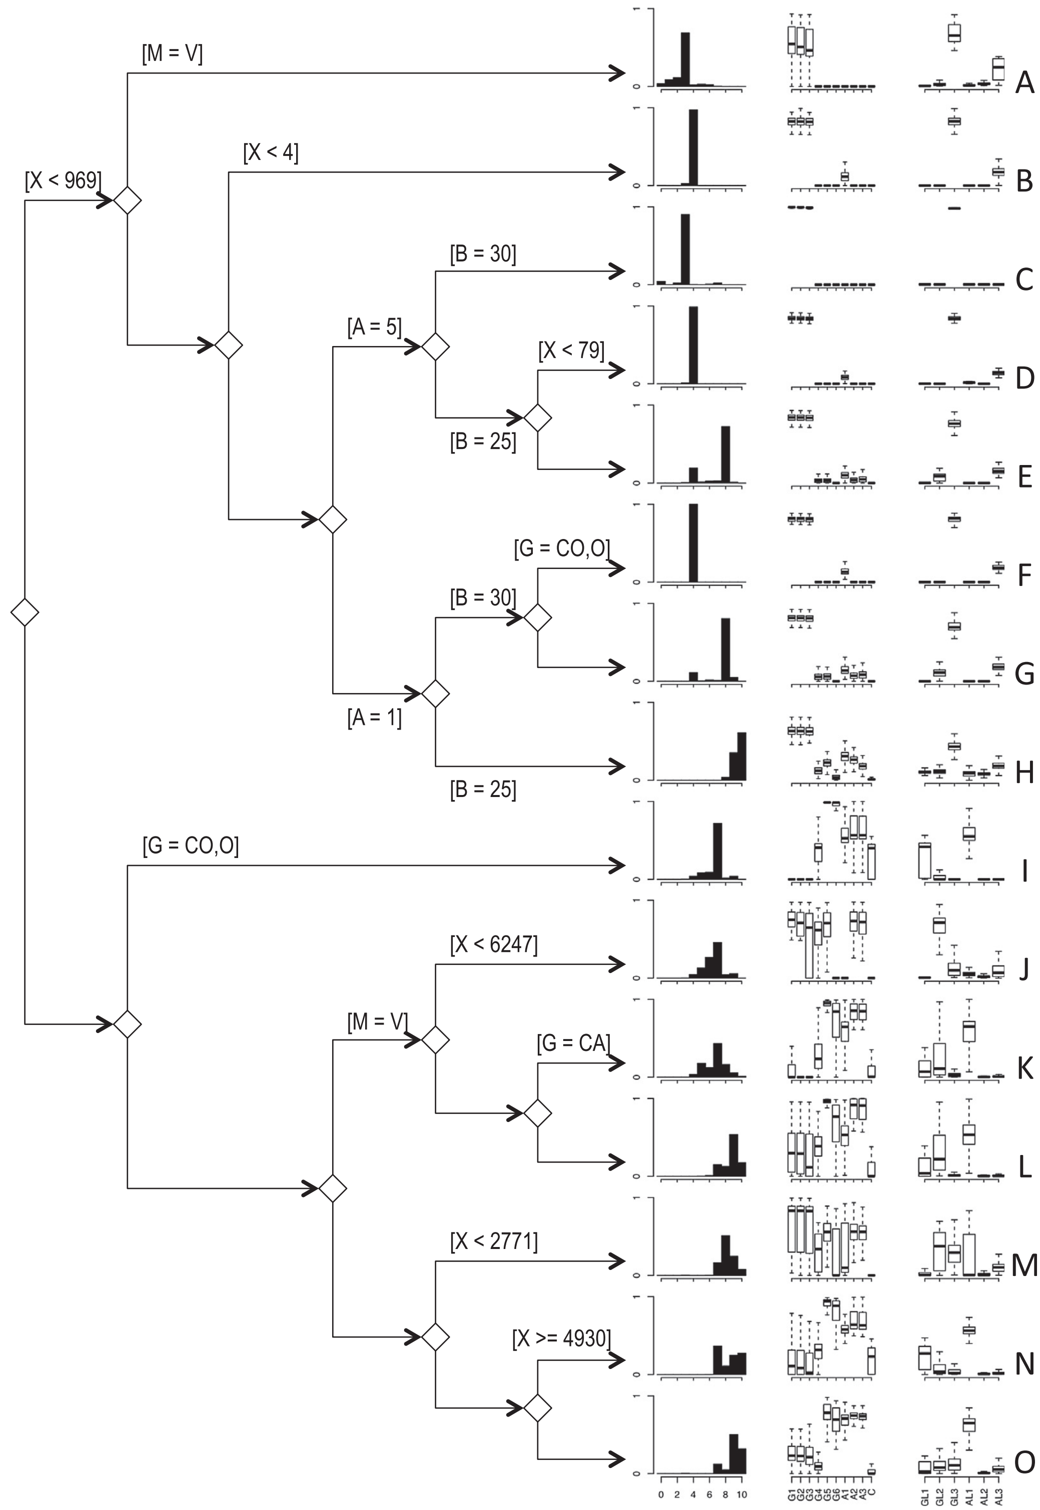

Fig. 5. Recursive partitioning, with richness as the response variable, and mean annual expenditure (X), government rule (G), market (M), break-even threshold (B) and aspiration (A) as explanatory variables. For each leaf node, three graphs are plotted. From left to right, these are: a histogram of richness and boxplots of species and land use occupancies. The species boxplot shows, from left to right, distributions of G1–6, A1–3 and C mean occupancies over the last 100 time steps of each run. The bottom and top of each box shows the first and third quartile, with a thicker line drawn to depict the median. ‘Whiskers’ (dashed lines) are drawn at the most extreme datapoints that are not more than 1.5 times the interquartile range from the box. Similarly, the land use boxplot shows distributions of mean occupancies of GL1–3 and AL1–3 from left to right.

land managers are satisficing on AL1, which for many is the initial land use their parcel is assigned, and hence do not experiment with GL2.

- Node N, meanwhile, shows GL1 being used more than GL2 or GL3, despite each GL1 parcel making a loss of 3 units of wealth. This is possible because land managers may own more than one parcel, and assess their aspirations at the enterprise level. A land manager having GL2 on one parcel can cross-subsidise GL1 on another and still satisfice in this scenario when the incentive is 10.5. When the incentive is 15.5, a land manager with AL1 can also cross-subsidise GL1. With GL1 being used more than GL2 or GL3, there is much less refuge for species G1–3 from species C, which outcompetes them. The result is more likely to be a richness of 7.

It is worth comparing the above analysis of A/F/25/5 with what would happen in this scenario if land managers optimised their profit, to see the effect this would have on nonlinearity. With zero incentive, GL3 has the highest profit (at 8, it is slightly higher than AL3's 7.5). Assuming the incentive for GL2 and AL1 is the same (as has been the case in these experiments), GL2 will always return 5 units of wealth more than AL1 in a fixed market. GL2 has higher profit than GL3 once the incentive is more than 5.5. At incentives below 5.5, the landscape would be entirely GL3, and only three species (G1–3) would survive; above 5.5, the landscape would be entirely GL2, providing habitats for all species except A1, though poor habitats for G6 and C. Landscape scale richness could be in the range 4 (if G6 died out, and C outcompeted G1–3 and then died out) to 9 (if G6 and C survived, the latter in low enough numbers that it did not outcompete G1–3). Neither extreme is particularly likely given the low level of habitat for G6 and C, and a richness of 7 may reasonably be expected to be the outcome.

Increasing the break-even threshold from 25 to 30 (Fig. 4(b), top right graph), does not change the incentive point at which GL2 overtakes GL3 – the profit of both is simply reduced by 5 units of wealth. Thus, with optimising land managers, A/F/30/5 and A/F/25/5 would look exactly the same from a land use perspective: at incentive 5.5, the landscape will shift from GL3 to GL2. Aspiration is, of course, irrelevant to optimising algorithms, and hence A/F/25/1 (Fig. 4(b), top left graph) would also look the same as A/F/25/5.

For A/F/30/5, Fig. 4(b) shows a greater incentive required to achieve the highest level of richness, and no decline to richness 7 by incentive 50, as there is in A/F/25/5. Ignoring the latter point briefly, it is reasonable to postulate that the effect of increasing break-even threshold has been to increase the amount of incentive required to achieve the same effect. This is supported by the vertical lines in Fig. 4 representing break-points drawn by a classification tree with incentive as the explanatory variable and richness as the response. Turning back to the difference between A/F/25/5 and A/F/30/5 at incentive 50, the incentive 100 runs were excluded because mean annual expenditure was too high (34,500, averaging about 73% of income). In A/F/30/5, all these runs have richness 7, just as the A/F/25/5 runs do at incentive 50.

Whilst A/F/25/1 shows a decline in richness from 10 to 7 as expenditure increases from 0 to 30, from node H the lower break-even and aspiration thresholds mean that even at low expenditure, though GL3 and AL3 dominate, they can more easily cross-subsidise other land uses enabling a wider diversity of species to survive. Hence, A/F/25/1 is sufficiently benign that the behaviour observed at low incentives in A/F/25/5 and A/F/30/5 does not occur.

The effects of outcome and clustered incentives are less straightforward to compute from an optimising perspective, as income depends on the behaviour of others. For clustered incentives, only the behaviour of the immediate neighbours is important; for outcome incentives, all other agents' decisions affect species' occupancy of the landscape, though arguably those nearer are more important than those further away. We may consider two extremes – one in which all farmers 'defect' by choosing the land use providing the highest profit assuming their neighbours do not attempt to acquire incentives, and another in which all farmers 'co-operate' by assuming their neighbours will attempt to acquire incentives.

In scenario CA/F/25/5, the 'defect' extreme will have the same threshold as A/F/25/5; the difference will be that when all managers switch to GL2 from GL3 at incentive 5.5, they will receive nine times more money for the same incentive value. The 'co-operate' extreme will lower the incentive threshold – an incentive of 11/18 would be enough to cause these optimising managers to switch from GL2 to GL3. The graph in the bottom right of Fig. 4(b) has a similar shape to those of A/F/25/5 and A/F/30/5, but with changes in outcome occurring at lower incentive values.

Scenario O/F/25/5 has the further complication that it is influenced by species' occupancy probabilities. In the 'defect' case, the fact that the incentive is outcome based means that this complication is irrelevant if we assume that one parcel is not sufficient to sustain species occupancy: land managers will never choose GL2 no matter what the incentive, and since GL3 provides habitat for G3, they will still receive some incentive. Turning to the 'co-operate' case, GL2 offers habitat for all the awardable species – at an 80% level for A2, A3, G3 and G5, but only at a 20% level for species G6. AL1 offers habitat at a 100% level for A2, A3, G5 and G6. For optimisers, using $P_i(x)$ to represent the agents' perceived probability of occupancy of species x on land use i, the choice will be between the more certain profit of $8 + \text{Incentive} \times P_{\text{GL3}}(\text{G3})$ when using GL3, and a riskier profit with GL2 of $2.5 + \text{Incentive} \times (P_{\text{GL2}}(\text{A2}) + P_{\text{GL2}}(\text{A3}) + P_{\text{GL2}}(\text{G3}) + P_{\text{GL2}}(\text{G5}) + P_{\text{GL2}}(\text{G6}))$. Using $\Delta P_{\text{GL2, GL3}}$ to represent the difference in probabilities of occupancy of awardable species between GL2 and GL3, the threshold between GL3 and GL2 adoption is when $\text{Incentive} > 5.5/\Delta P_{\text{GL2, GL3}}$. If the differences between the habitat provision of AL1 and GL2 are sufficient that $\Delta P_{\text{AL1, GL2}}$ is positive, then assuming $\Delta P_{\text{AL1, GL2}} << \Delta P_{\text{GL2, GL3}}$ there will be a second threshold at which AL1 is adopted over GL2 when $\text{Incentive} > 5/\Delta P_{\text{AL1, GL2}}$. AL1 supports fewer species than GL2, and hence the higher incentive will result in lower landscape-scale species richness.

The fact that there is a social dilemma in the clustered and/or outcome incentive rules raises the interesting possibility of bifurcation in emergent outcomes, as has been demonstrated, for example, in the prisoner's dilemma (e.g. Perc, 2006). Bifurcation is a technical term associated with the nonlinear dynamics literature, and without further analysis should perhaps be used with caution here. Nevertheless, Fig. 4(b) shows for scenario O/F/25/5 that with incentive 1 roughly half the runs (31) have richness 4, and a similar number of runs (25) have richness 8, suggesting that with these parameters, stochasticity in either the initial conditions or the ensuing dynamics has resulted in notably different biodiversity outcomes. With respect to the foregoing discussion, land managers' choices are of greater relevance than species richness, however. Since the question of whether land managers 'co-operate' or 'defect' is not specifically addressed by the decision-making algorithm, this matter may best be analysed through further research.

A variable market also complicates matters. From Fig. 2, none of the land uses satisfices or breaks even all of the time. (The minimum gross income from GL3, the land use with the highest minimum gross income, is 24.) The difference in 'phase' between the GLx land use market and ALx market means that optimising land managers will change land use. If all land managers are optimisers, the species that can possibly survive an optimising population of land managers will be those that can survive in both GHx and AHx habitats: A2, A3, G5 and G6. The habitats for these species are provided by GL1, GL2, AL1 and AL2. Hence, at low incentive values, whatever the government rule, optimising land managers switching between GL3 and AL3 will cause the extinction of all species. In scenario A/V/25/5, only GL2 and AL1 are incentivised, though these provide some habitat for all four of the above species, GL2 only has 20% habitat for G6, meaning that its long-term survival is unlikely. Thus a richness of 3 should be the expected outcome from optimising land managers in this scenario, once the incentive is high enough that managers never choose GL3 or AL3 over GL2 or AL1. This means the ALx market is the more critical, as the difference in yield between AL1 and AL3 is larger than that between GL3 and GL2, and hence the threshold in incentive between a richness of 0 and a richness of 3 is when the maximum profit for AL1 is more than the maximum for AL3: 13.5. A/V/25/5 is shown in the bottom left of Fig. 4(d). Interestingly, it is similar to A/F/25/5, except that it does not diminish to a richness of 7 by incentive 50, but stays around 9. From node L in Fig. 5, higher expenditure in A/V/25/5 is associated with GL2 and AL1 being used, which can be accounted for through land managers experimenting when AL1 is unprofitable, and learning that GL2 provides a higher rate of return than GL3.

Comparing Fig. 4(a) and (b) with (c) and (d), it is clear that with a variable market and outcome-based incentive rule instead of a fixed market and activity-based incentive rule, the other scenario variables do not have the same effect on the relationship between incentive and richness.

From the above, two important points emerge. First, a nonlinear relationship between incentive and richness arises when land managers optimise their profits, with thresholds associated with incentive levels that compensate for more profit-making land uses. Second, while the agents' non-profit-maximising decision algorithm means the thresholds are less pronounced, it also causes further nonlinear behaviour in the incentive/richness relationship. A corollary of the latter point is that non-profit-maximising behaviour can drive spatial variation in habitat beneficial to species, separately from any biophysical considerations (which have not been addressed in these experiments). The following lists the effects that arise from the decision algorithm used that are apparent from the foregoing analysis:

- Satisficing means that, no matter how high the incentive for other land uses, any land use that meets aspirations will be chosen. This will cause greater landscape diversity than optimisation, particularly if aspirations and break-even thresholds are low. (Nodes I, N and O in Fig. 5, featuring AL1 dominating instead of GL2, show this happening.)
- Assessing aspirations at the enterprise rather than land parcel level allows a land use making enough profit to cross-subsidise another that makes a loss. (Nodes I, K, L, N and O show use of GL1.)
- In variable markets, the relationship between the length of time for which a land parcel fails to meet aspirations and a land manager's change delay parameter will be important. Longer change delays can mean land uses survive that fail to satisfice for shorter periods, preserving their associated habitats. (Nodes A and J, for example, show survival of G1–3, when switching between the ALx and GLx land uses through optimised decision-making would cause their extinction.)

4. Discussion

The results show that the relationship between incentive and biodiversity is far from trivial. A naïve expectation would be that more money buys more biodiversity. The relationship, however, is not linear. There seem to exist thresholds of expenditure (Figs. 4 and 5), which relate to habitat provision, below which landscape-scale biodiversity remains low (i.e. many species extinctions occur, because their habitat is not sufficiently common). Whilst intervening samples on the incentive axis would give a clearer picture of the shape of the nonlinearities (specifically, whether they are discontinuities, such as in the analysis of the optimal case for A/F/25/5 above), our use of the term 'threshold' is consistent with other authors' (e.g. Samhouri et al., 2010). Spending money below such thresholds thus seems ineffective if the objective is saving species from extinction, and finding such thresholds is therefore important. It is not clear that fixed incentives or even mechanisms such as auctions would be able to find such thresholds, and real-world incentive schemes in any case tend to be aimed at the individual farmer rather than the landscape scale. As such they might be efficient to implement in the short term, but are probably ineffective in the long term at maintaining landscape-scale species diversity. (Exploring these issues is the possible subject of future work.) A more sophisticated model would also recognise that there are limits to biodiversity, suggesting a sigmoidal shape to the relationship.

Of particular interest are those results such as in Fig. 4(a) and (b) showing biodiversity reaching a peak at a point before maximum incentive/expenditure is reached – these suggest that biodiversity is not necessarily a monotone increasing function of incentive. Comparison with profit maximisation shows that various aspects of the land managers' decision-making algorithm create situations where lower incentives can outperform higher incentives with respect to species richness, though in some scenarios (e.g. O/F/25/5, discussed above), this can occur with profit-maximising agents if species' occupancy probabilities have appropriate values. Whether these aspects of decision-making are artefactual or not, it is clear from our results that the success of an agri-environmental incentive scheme is sensitive to the way farmers make decisions in subtle ways that would be difficult, if not impossible, for a government to cater for in the real world.

What is also clear from the results is that there is not a *general* pattern to the incentive/richness relationship, highlighting the importance of the context in which this relationship occurs. Not only do each of the non-government-related scenario variables have an effect on the richness, these effects interact with each other in a context-sensitive way: certain features of land manager decision-making only become apparent when, for example, aspirations are too high, or the variable market is used. Edmonds and Akman (2002), in a special issue on context, contrast approaches in Artificial Intelligence to context treating it as a series of properties of the environment the inclusion of which broadens generality, with their stance, which is based on a "messy and contingent" (p. 234) worldview. They argue for a more cautious approach to generality, which starts with a search for aspects of context that are shared at a local (spatial/temporal/organisational/cultural) scale.

Ostrom et al. (2007) introduce a series of articles challenging the views that general solutions (panaceas) to issues of environmental degradation in socio-ecological systems exist and can be deduced from simple, predictive models of them. Ostrom (2007) proposes instead a framework for analysing socio-ecological systems that is aimed at developing answers to questions on: (a) the relationship between institutions for resource distribution and governance and patterns of social and environmental outcomes; (b) the possible bottom-up development of such institutions without other external incentives; and (c) the robustness of various institutional arrangements to shocks. The framework comprises a number of variables that empirical studies of socio-ecological systems could report on to facilitate meta-analyses of case studies. Whilst she acknowledges that these variables are neither exhaustive nor universally relevant, Edmonds (2012a) points out that in any complex system, there are no bounds to the possible causes of an event, suggesting a potentially limitless number of variables. For example, biodynamic farmers take into consideration the positions of celestial bodies in land management practices; it is doubtful that simulations of socio-ecological systems will ever include a sub-model simulating the positions of stars and planets in the night sky.

The issue of context-sensitivity and generality of governance of socio-ecosystems raises questions of the appropriate scale at which to design and implement agri-environmental incentive schemes.

Scottish biodiversity incentivisation is largely embedded in the Scottish Rural Development Programme[2] – SRDP. This scheme includes rural priorities: the relative importance of five outcomes (not all of which have biodiversity or more general pro-environmental potential benefits) have been agreed with stakeholders in each of 11 regions of Scotland. The context sensitivity of our results lends theoretical support to a more localised element to incentive schemes, However, there is no reason to believe that a regional approach is optimal with respect to locality. Indeed, Brondizio et al. (2009) argue that there is no fixed appropriate scale for ecosystem governance. Until recently, the Farming and Wildlife Advisory Group (FWAG) offered advice to farmers in the UK, tailored to their business type, location, aspirations and financial resources. Insofar as advice tailored to location factored in a co-ordinated approach to conservation at the landscape scale, the recent closure of this service (Driver, 2011) may represent a significant lost opportunity for the integrated cross-scale management of biodiversity.

Findings from the conservation literature have already highlighted that if grazing intensity is too low, there can be a negative impact on biodiversity (Wallis de Vries et al., 1998) and this is reflected in the setup for these simulations through the lowest intensity grazing land use (GL1) providing habitat for a competitor species (C) that outcompetes species G1, G2 and G3. However, in these simulations, the government agent is assumed to know about this, and provides incentives for GL2 rather than GL1 in the activity rules, and for G3 in the outcome rules. What these results show is that despite providing incentives to manage landscape-scale biodiversity, the response of biodiversity to government policy can be non-monotonic as a result of land managers' decision-making processes. In the model, satisficing and the scale at which aspirations are assessed by land manager agents are key contributors to higher incentives providing lower biodiversity, as these lead to AL1 providing enough income to satisfice and cross-subsidise use of GL1 in the same enterprise, and hence no motivation to try GL2, which provides much poorer habitat for C. Nevertheless, the analysis of scenario O/F/25/5 above shows that biodiversity loss at higher incentive values is possible if the model used profit-maximising land use selection algorithms.

Though it would be a mistake to regard farmers as less interested in profit than other businesses, profit maximisation is an unrealistic goal, particularly given costs that may be associated with some changes in commodity or land management regime. Laboratory experiments comparing resource allocation of human actors with utility-maximising agents suggest that non-maximising decision-making algorithms should be deployed in agent-based models of land use change (Evans et al., 2006). Furthermore, qualitative research of Scottish farmers has revealed that farmer identity can play a significant role in determining land use activity, and in particular, as shown by Burton (2004), a predominantly productivist mindset means that conservation activities are seen as 'not farming'. A more recent study of farmers in the Upper Deeside area of northeast Scotland identified three motivations for incentive scheme uptake: the scheme offers payments for activities they are already doing; a sense of entitlement to farming subsidies; and an opportunistic culture (Sutherland, 2010). However, Sutherland's (2011) study of English farmers finds that there are a number of farmers who regard themselves as 'effectively' or 'semi-' organic, owing to their relatively low use of chemical inputs, but without organic certification or agri-environmental incentives to do so. Though this may be driven in part by legislation on chemical inputs, and by increasing input costs, she found a more positive attitude towards some organic ideals among the farmers she interviewed, leading her to borrow the term 'organification' from Rosin and Campbell (2009) to describe trends in conventional farming.

Though the results show that government policy, both the way it is administered and the amount spent, can have an effect on biodiversity, these studies have all kept both the administration and amount constant throughout any one run. Thus the results cannot necessarily be used to infer that a particular change in policy (either increasing incentive or changing administration rule) will achieve the particular effect on biodiversity suggested here *within a single run*. In separate work, we have explored the possibilities of using techniques from control theory (Nise, 2004) to navigate the space of possible incentive schemes in an initial investigation (Polhill et al., 2010a).

Most incentives schemes reported in the literature are based on agreed activities. These entail cost-sharing or total compensation to carry out conservation-related activities (e.g. Hacker et al., 2010), or to refrain from some practices, such as excessive grazing or mowing grass before a given date or with some frequency (e.g. Humbert et al., 2009; Berg and Gustafson, 2007).

Points-based schemes are an alternative that have been used to reward land use activities in various countries from Latin America (Pagiola et al., 2004) to Scotland. However, doubts have been raised in many studies regarding the effectiveness of existing incentive schemes (Stoate et al., 2009; Goldman et al., 2007; Vickery et al., 2002, 2004). This is partly due to the limited and fragmented scale at which they often operate, when what is really needed is land management action at the landscape scale (e.g. Parkhurst and Shogren, 2007; Davies et al., 2004; Selman, 2006; Gimona and Van der Horst, 2007; Pelosi et al., 2010). In our model results, given the right incentives and land managers' circumstances, often enough land managers took on incentives to promote the landscape-scale species diversity. However, our model does not include productivist incentives, which often compete with environmental ones and might contribute to explaining the poor performance of real incentives schemes lamented by the authors above. The SRDP, for example, provides farmers with a series of options of activities for which they can obtain incentives, not all of which have a biodiversity benefit. Moreover, non-compliance is a possible land manager option in the real world that is not simulated here, and this requires potentially costly monitoring. These problems are therefore similar to those associated with a regulation-based approach.

Another alternative increasingly being discussed is rewarding based on outcomes, which aims to overcome the risk of cheating and the asymmetry of knowledge between land manager and government, and has proven successful in number of cases. These include a range of ecosystem services in the USA, and in different areas of Australia (Latacz-Lohmann and Schilizzi, 2005). However, while there are theoretical reasons that make economists favour such schemes (e.g. Horowitz et al., 2009), evidence of their actual superiority in terms of delivering conservation is not clear-cut, due to scarcity of evaluation studies. Our results suggest outcome-based schemes may be less sensitive to other scenario variables with possible influences on behaviour (node I in Fig. 5). An important disadvantage associated with many outcome-based schemes is the time delay between activities and outcomes, which means that most of the risk is borne by land managers. This might make such schemes less attractive. For this reason, some schemes combine more traditional conservation activity incentives with further outcome-based rewards if the activities prove effective (Whitten et al., 2008).

Sutherland et al. (2012) also note environmental benefits deriving from clustering of organic farms, and we have explored the use of a clustering rule to incentive schemes in these experiments. Our results suggest that it can have an effect only in a restricted set

[2] http://www.scotland.gov.uk/Topics/farmingrural/SRDP.

of contexts (nodes K and L in Fig. 4), and one that is surprisingly slightly detrimental (at least in terms of the modal species richness of the applicable runs). Our earlier results with clustered incentives found they created a more stable habitat for vulnerable species (A3 and G6) than non-clustered incentives, though at a lower overall level of occupancy that nevertheless enabled longer persistence times (Gimona and Polhill, 2011). The results here suggest that this difference between clustered and non-clustered incentivisation might only be observed at higher levels of mean annual expenditure. However, there are various ways that a clustering rule could be implemented, including the potentially interesting prospect of simulating land manager agents explicitly arranging to co-operate in applying or bidding for participation in such schemes. Clustering is also only one of several ways in which an incentive scheme could be designed to deliver landscape-rather than farm-scale biodiversity.

A proper comparison of the cost-effectiveness of incentive schemes would also take into account transaction cost, due to assuring compliance. Schemes with outcomes that can be easily monitored (e.g. through remote sensing or rapid assessment methods) are clearly more efficient than those that cannot. In many real-world cases, however, considerations about cost-effectiveness, including the effect of transaction costs, are under-examined (Coggan et al., 2010) and therefore secondary with respect to biophysical ones. The government agent in our model is omniscient with respect to species and activity on each patch of land, and these costs are not factored in to the model. Future work could examine the effects of different sampling mechanisms for monitoring species presence.

Edmonds's (2012b) more recent reflections on context in social simulation include criticisms of the general application of models to different contexts. Whilst these are, perhaps, more focused on application of models of magnetism to decision-making in social systems (e.g. Sznajd-Weron and Sznajd, 2000), it is reasonable to question whether the decision-making algorithm used here is generally applicable across all the scenarios studied. For example, in the variable market case, perhaps managers should have a price prediction submodel; or in the outcome incentive case, managers should attempt to develop predictive models of species occupancy based on their own and neighbouring activities. In cases of very high incentives, the satisficing heuristic may need adjustment. For example, perhaps there is a 'temptation' threshold, whereby, although managers are content with their present income, with high enough incentives, they will nevertheless change activities in order to obtain them. An alternative would be to adjust aspirations – either on the basis of incentives themselves, or on evidence of higher profits made by neighbours. The question is whether making such adjustments would detract from our observations – would the nonlinearity, dependency on non-optimality of decision-making, and context sensitivity go away? This is the potential subject of future research. An alternative would be to *assume* that our observations would go away if the proposed modifications were made – this would make policy design easier. There is a precedent for making convenient 'heroic assumptions' (to borrow a term from Johnson (1998)), such as that heterogeneities among agents cancel out (Weisskopf, 1955); but suspicion of these is among the motivations for agent-based modelling in the first place.

5. Conclusion

We have shown that there are turning points in biodiversity delivered for incentive invested, that such delivery is not a monotone increasing function of incentive, and that, although these turning points depend partly on incentive scheme delivery, they also depend on other aspects of the ecological, economic, psychological and social agricultural environment that would be very difficult for a government to ascertain. As such, these simulations lend theoretical support to incentive schemes that feature components tailored to more local scales, perhaps even down to the farm gate itself.

Acknowledgements

This work was funded by the Scottish Government Rural Affairs and the Environment Portfolio Strategic Research Theme 1 (Ecosystem Services). Contributions to FEARLUS–SPOMM software used in these experiments have been made by Aurore Catalan, Davy Decorps, Guillaume Dury, Luis Izquierdo, Alistair Law and Julien Oleon. We also acknowledge IT support from Bode Akitoye, and thank two anonymous referees for helpful comments on an earlier draft.

Appendix A

Fig. A.1 shows the scenarios not already shown in Fig. 4. An explanation of the plots is given in the caption to Fig. 4.

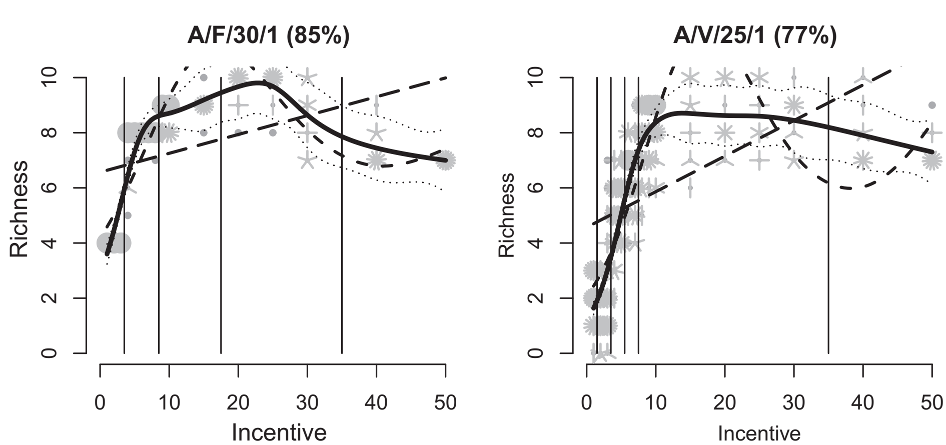

Fig. A.1. Plots in the style of Fig. 4 for the scenarios not depicted therein.

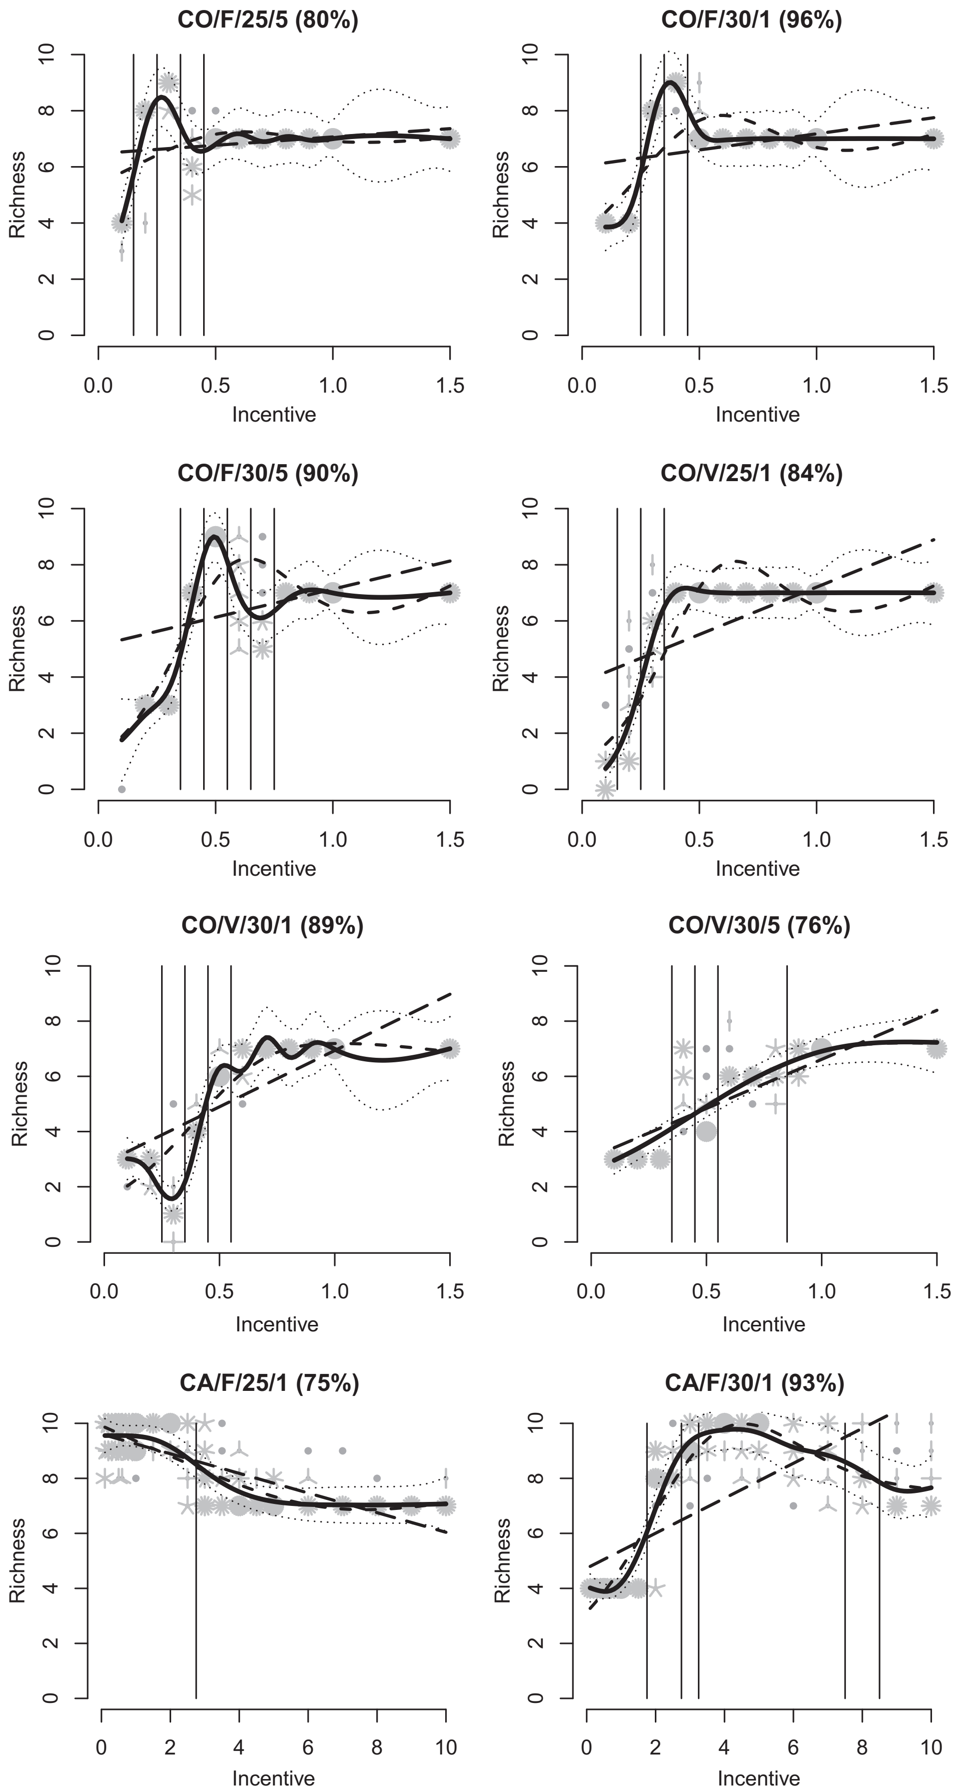

Fig. A.1. (*continued*).

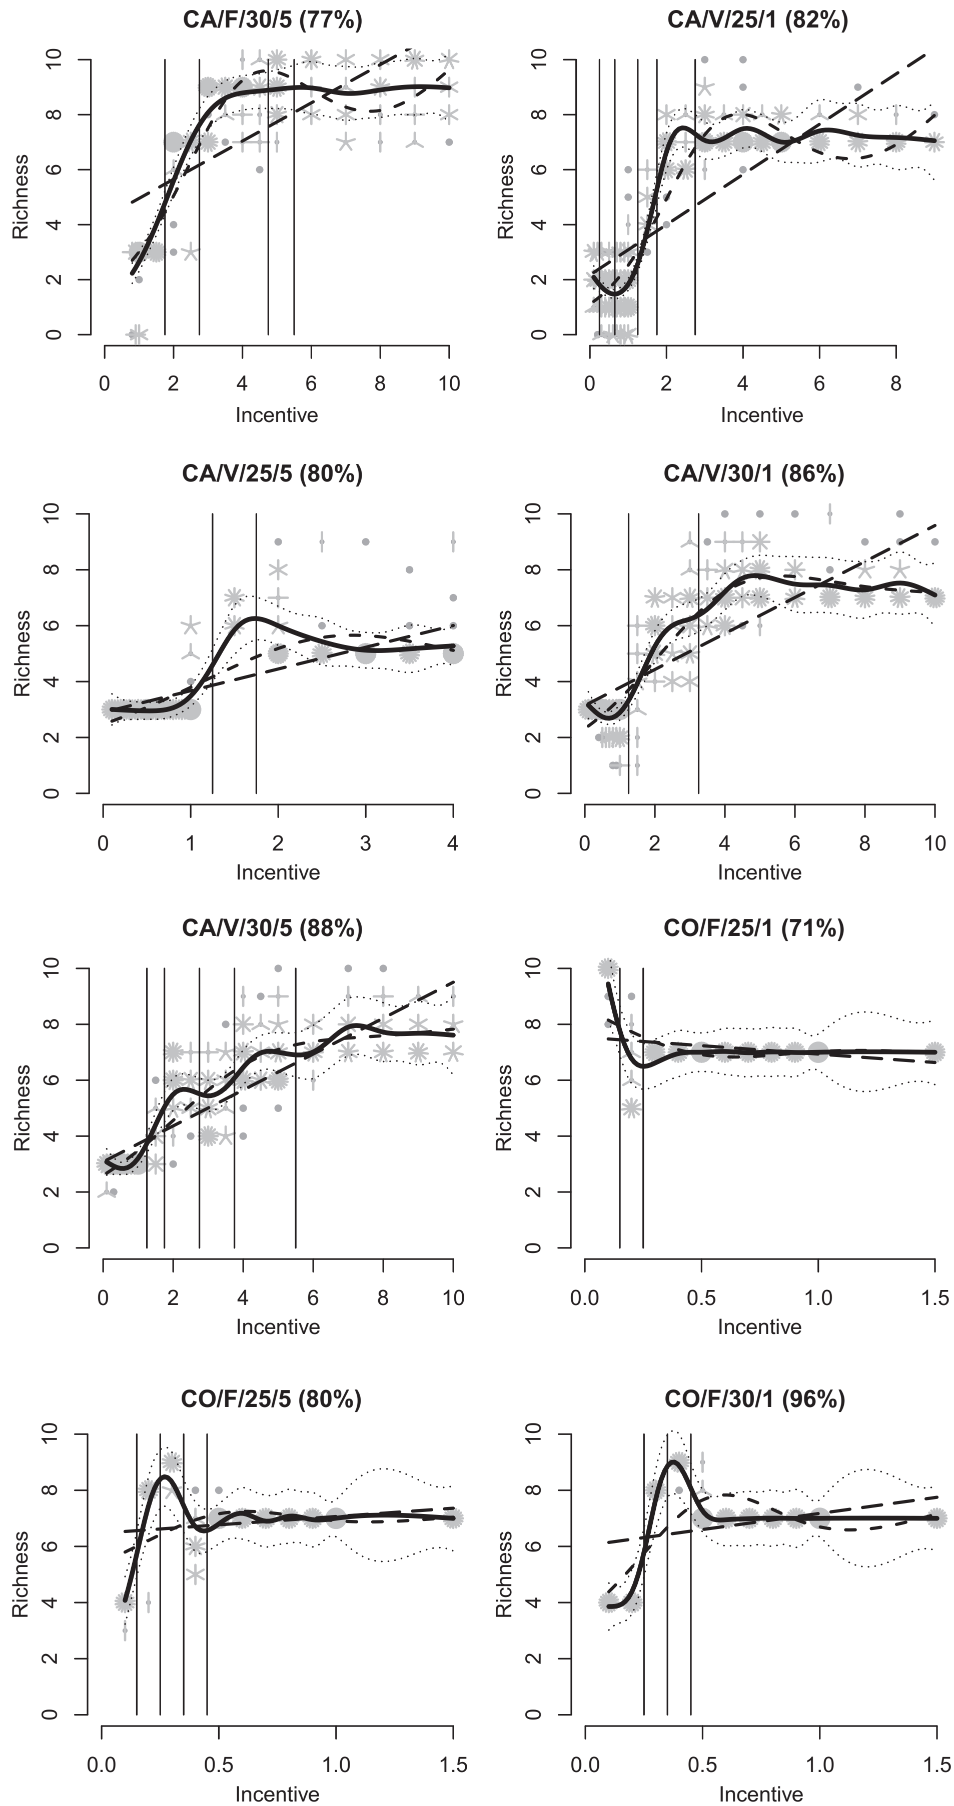

Fig. A.1. (*continued*).

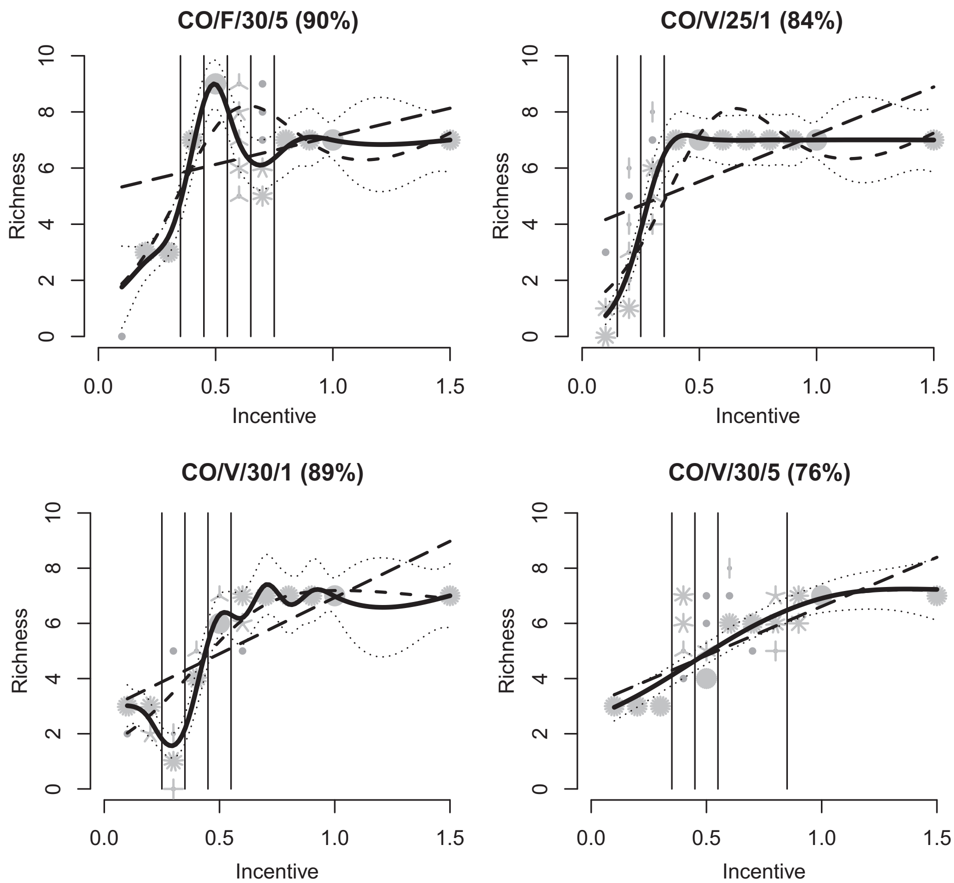

Fig. A.1. (*continued*).

References

Aamodt, A., Plaza, E., 1994. Case-based reasoning: foundational issues, methodological variations, and system approaches. AI Commun. 7 (1), 39–59.

Akaike, H., 1973. Information theory and an extension of the maximum likelihood principle. In: Petrov, B.N., Caski, F. (Eds.), Second International Symposium on Information Theory. Akademiai Kiado, Budapest, pp. 267–281.

Akaike, H., 1974. A new look at the statistical model identification. IEEE Trans. Automatic Control 19 (6), 716–723.

An, L., Linderman, M., Qi, J., Shortridge, A., Liu, J., 2005. Exploring complexity in a human-environment system: an agent-based spatial model for multidisciplinary and multiscale integration. Ann. Assoc. Am. Geogr. 95, 54–79.

Anselme, B., Bousquet, F., Lyet, A., Etienne, M., Fady, B., Le Page, C., 2010. Modelling of spatial dynamics and biodiversity conservation on Lure mountain (France). Environ. Model. Softw. 25, 1385–1398.

Axelrod, R., 1997. The Complexity of Cooperation: Agent-based Models of Competition and Collaboration. Princeton University Press.

Axelrod, R., Riolo, R.L., Cohen, M.D., 2002. Beyond geography: cooperation with persistent links in the absence of clustered neighborhoods. Pers. Soc. Psychol. Rev. 6, 341–346.

Banks, P.B., Norrdahl, K., Korpimäki, E., 2000. Nonlinearity in the predation risk of prey mobility. Proc. R. Soc. Lond. B 267, 1621–1625.

Benton, T.G., Vickery, J.A., Wilson, J.D., 2003. Farmland biodiversity: is habitat heterogeneity the key? Trends Ecol. Evol. 18, 182–188.

Berg, A., Gustafson, T., 2007. Meadow management and occurrence of corncrake *Crex crex*. Agric. Ecosyst. Env. 120, 139–144.

Boero, R., Castellani, M., Squazzoni, F., 2004. Cognitive identity and social reflexivity of the industrial district firms. Going beyond the 'complexity effect' with an agent-based computational prototype. In: Lindermann, G., Modt, D., Paolucci, M. (Eds.), Regulated Agent-based Social Systems. Springer Verlag, pp. 48–69.

Boero, R., Squazzoni, F., 2005. Does empirical embeddedness matter? Methodological issues on agent-based models for analytical social science. J. Artif. Soc. Soc. Simul. 8 (4), 6. http://jasss.soc.surrey.ac.uk/8/4/6.html.

Boulanger, P.-M., Bréchet, T., 2005. Models for policy-making in sustainable development: the state of the art and perspectives for research. Ecol. Econ. 55, 337–350.

Brady, M., Sahrbacher, C., Kellermann, K., Happe, K., 2012. An agent-based approach to modeling impacts of agricultural policy on land use, biodiversity and ecosystem services. Landscape Ecol. 27, 1363–1381.

Breiman, L., Friedman, J., Olshen, R.A., Stone, C.J., 1984. Classification and Regression Trees. Chapman & Hall/CRC, Boca Raton, Florida.

Brondizio, E.S., Ostrom, E., Young, O.R., 2009. Connectivity and the governance of multilevel social-ecological systems: the role of social capital. Annu. Rev. Environ. Resour. 34, 253–278.

Burton, R.J.F., 2004. Seeing through the 'good farmer's' eyes: towards developing an understanding of the symbolic value of 'productivist' behaviour. Sociol. Rural. 44, 195–215.

Carvalheiro, L.G., Veldtman, R., Shenkute, A.G., Tesfay, G.B., Pirk, C.W.W., Donaldson, J.S., Nicolson, S.W., 2011. Natural and within-farmland biodiversity enhances crop productivity. Ecol. Lett. 14, 251–259.

Coggan, A., Whitten, S.M., Bennett, J., 2010. Influences of transaction costs in environmental policy. Ecol. Econ. 69, 1777–1784.

Connell, J.H., 1978. Diversity in tropical rain forests and coral reefs. Science 199, 1302–1309.

Davies, B., Blackstock, K., Brown, K., Shannon, P., 2004. Challenges in Creating Local Agri-environmental Cooperation Action Amongst Farmers and Other Stakeholders. The Macaulay Land Use Research Institute, Report for Scottish Executive Environmental and Rural Affairs Department. Aberdeen.

Dean, S.J., Gumerman, G.J., Epstein, J.M., Axtell, R., Swedlund, A.C., Parker, M.T., McCarroll, S., 2000. Understanding Anasazi culter change through agent-based modeling. In: Kohler, T.A., Gumerman, G.J. (Eds.), Dynamics in Human and Primate Societies: Agent-based Modeling of Social and Spatial Processes. Oxford University Press, pp. 179–205.

Driver, A., 2 November 2011. FWAG to make announcement on its future. Farmers Guard.

Edmonds, B., October 2012a. Complexity and context-dependency. Found. Sci., http://dx.doi.org/10.1007/s10699-012-9303-x.

Edmonds, B., 2012b. Context in social simulation: why it can't be wished away. Comput. Math. Organ. Theory 18 (1), 5–21.

Edmonds, B., Akman, V., 2002. Editorial: context in context. Found. Sci. 7 (3), 233–238.

Evans, T., Sun, W., Kelley, H., 2006. Spatially explicit experiments for the exploration of land-use decision-making dynamics. Int. J. Geog. Inf. Sci. 20 (9), 1013–1037.

Gilpin, M.E., Hanski, I.A., 1991. Metapopulation Dynamics: Empirical and Theoretical Investigations. Academic Press.
Gimona, A., Polhill, J.G., 2011. Exploring robustness of biodiversity policy with a coupled metacommunity and agent-based model. J. Land Use Sci. 6 (2–3), 175–193.
Gimona, A., Polhill, G., Davies, B., 2011. Sinks, sustainability, and conservation incentives. In: Liu, J., Hull, V., Morzillo, A., Wiens, J. (Eds.), Sources, Sinks and Sustainability. Cambridge University Press, Cambridge U.K. (Chapter 8).
Gimona, A., Van der Horst, D., 2007. Mapping hotspots of multiple landscape functions: a case study on farmland afforestation in Scotland. Landsc. Ecol. 22, 1255–1264.
Goldman, R.L., Thompson, B.H., Daily, G.C., 2007. Institutional incentives for managing the landscape: inducing cooperation for the production of ecosystem services. Ecol. Econ. 64, 333–343.
Gotts, N.M., Polhill, J.G., 2009. When and how to imitate your neighbours: lessons from and for FEARLUS. J. Artif. Soc. Soc. Simul. 12 (3), 2. http://jasss.soc.surrey.ac.uk/12/3/2.html.
Gotts, N.M., Polhill, J.G., Law, A.N.R., 2003. Aspiration levels in a land use simulation. Cybern. Syst. 34, 663–683.
Grimm, V., Berger, U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J., Goss-Custard, J., Grand, T., Heinz, S.K., Huse, G., Huth, A., Jepsen, J.U., Jørgensen, C., Mooij, W.M., Müller, B., Pe'er, G., Piou, C., Railsback, S.F., Robbins, A.M., Robbins, M.M., Rossmanith, E., Rüger, N., Strand, E., Souissi, S., Stillman, R.A., Vabø, R., Visser, U., DeAngelis, D.L., 2006. A standard protocol for describing individual-based and agent-based models. Ecol. Model. 198, 115–126.
Grimm, V., Berger, U., DeAngelis, D.L., Polhill, J.G., Giske, J., Railsback, S.F., 2010. The ODD protocol: a review and first update. Ecol. Model. 221, 2760–2768.
Guzy, M.R., Smith, C.L., Bolte, J.P., Hulse, D.W., Gregory, S.V., 2008. Policy research using agent-based modeling to assess future impacts of urban expansion into farmlands and forests. Ecol. Soc. 13 (1), 37. http://www.ecologyandsociety.org/vol13/iss1/art37/.
Hacker, R.B., Jessop, P.J., Smith, W.J., Melville, G.J., 2010. A ground cover-based incentive approach to enhancing resilience in rangelands viewed as complex adaptive systems. Rangel. J. 32, 283–291.
Hald, A.B., 1999. The impact of changing the season in which cereals are sown on the diversity of the weed flora in rotational fields in Denmark. J. Appl. Ecol. 36, 24–32.
Hare, M., Deadman, P., 2004. Further towards a taxonomy of agent-based simulation models in environmental management. Math. Comput. Simul. 64, 25–40.
Hastie, T.J., Tibshirani, R.J., 1990. Generalized Additive Models. Chapman & Hall/CRC, Boca Raton, Florida.
Holyoak, M., Leibold, M.A., Hold, R.D., 2005. Metacommunities: Spatial Dynamics and Ecological Communities. University of Chicago Press.
Horowitz, J.K., Lynch, L., Stocking, A., 2009. Competition-based environmental policy: an analysis of farmland preservation in Maryland. Land Econ. 85, 555–576.
Humbert, J.Y., Ghazoul, J., Walter, T., 2009. Meadow harvesting techniques and their impacts on field fauna. Agric. Ecosyst. Env. 130, 1–8.
Johnson, P.E., 1998. Rational Actors Versus Adaptive Agents: Social Science Implications. 1998 Annual Meeting of the American Political Science Association, Boston.
Kwaiser, K.S., Hendrix, S.D., 2008. Diversity and abundance of bees (Hymenoptera: apiformes) in native and ruderal grasslands of agriculturally dominated landscapes. Agric. Ecosyst. Environ. 124, 200–204.
Lagabrielle, E., Botta, A., Daré, W., David, D., Aubert, S., Fabricius, C., 2010. Modelling with stakeholders to integrate biodiversity into land-use planning – lessons learned in Réunion Island (Western Indian Ocean). Environ. Model. Softw. 25, 1413–1427.
Latacz-Lohmann, U., Schilizzi, S. 2005. Auctions for Conservation Contracts: a Review of the Theoretical and Empirical Literature. A Report to the Scottish Executive Environmental and Rural Affairs Department, 15 October 2005. http://www.scotland.gov.uk/Resource/Doc/93853/0022574.pdf.
Lansing, J.S., Kremer, J.N., 1994. Emergent properties of Balinese water temple networks: coadaptation on a rugged fitness landscape. In: Langton, C.G. (Ed.), Artificial Life III. Addison-Wesley, pp. 201–203.
Lempert, R., 2002. Agent-based modeling as organizational and public policy simulators. Proc. Natl. Acad. Sci. U. S. A. 99, 7195–7196.
Levins, R., 1968. Evolution in Changing Environments. Princeton University Press.
Manson, S.M., 2005. Agent-based modeling and genetic programming for modeling land change in the Southern Yucatán Peninsular Region of Mexico. Agric. Ecosyst. Env. 111, 47–62.
Matthews, R.B., Polhill, J.G., Gilbert, N., Roach, A., December 2005. Integrating agent-based social models and biophysical models. In: Zerger, A., Argent, R.M. (Eds.), MODSIM 2005: International Congress on Modelling and Simulation. Modelling and Simulation Society of Australia and New Zealand, pp. 1617–1623. http://www.mssanz.org.au/modsim05/papers/matthews_rb1.pdf.
Meehan, T.D., Hurlbert, A.H., Gratton, C., 2010. Bird communities in future bioenergy landscapes of the Upper Midwest. Proc. Natl. Acad. Sci. U. S. A. 107, 18533–18538.
Moilanen, A., 1999. Patch occupancy models of metapopulation dynamics: efficient parameter estimation using implicit statistical inference. Ecology 80, 1031–1043.
Moilanen, A., 2004. SPOMSIM: software for stochastic patch occupancy models of metapopulation dynamics. Ecol. Model. 179, 533–550.
Naeem, S., Thompson, L.J., Lawler, S.P., Lawton, J.H., Woodfin, R.M., 1995. Empirical evidence that declining species diversity may alter the performance of terrestrial ecosystems. Philos. Trans. R. Soc. Lond. Ser. B 347, 249–262.
Nise, N.S., 2004. Control Systems Engineering. John Wiley, New York.
Ostrom, E., 2007. A diagnostic approach for going beyond panaceas. Proc. Natl. Acad. Sci. U. S. A. 104 (39), 15181–15187.
Ostrom, E., Janssen, M.A., Anderies, J.M., 2007. Going beyond panaceas. Proc. Natl. Acad. Sci. U. S. A. 104 (39), 15176–15178.
Pagiola, S., Agostini, P., Gobbi, J., de Haan, C., Ibrahim, M., Murgueitio, E., Ramírez, E., Rosales, M., Ruíz, J.P., 2004. Paying for Biodiversity Conservation Services in Agricultural Landscapes. Environment Department Paper No. 96. World Bank, Washington.
Parker, D.C., Hessl, A., Davis, S.C., 2008. Complexity, land-use modeling, and the human dimension: fundamental challenges for mapping unknown outcome spaces. Geoforum 39, 789–804.
Parkhurst, G.M., Shogren, J., 2007. Spatial incentives to coordinate contiguous habitat. Ecol. Econ. 64, 344–355.
Pelosi, C., Goulard, M., Balent, G., 2010. The spatial scale mismatch between ecological processes and agricultural management: do difficulties come from underlying theoretical frameworks? Agric. Ecosyst. Env. 139, 455–462.
Perc, M., 2006. Coherence resonance in a spatial prisoner's dilemma game. N. J. Phys. 8, 22.
Polasky, S., Doremus, H., Rettig, B., 1997. Endangered species conservation on private land. Contemp. Econ. Policy 15, 66–76.
Polhill, J.G., Gotts, N.M., Law, A.N.R., 2001. Imitative versus nonimitative strategies in a land-use simulation. Cybern. Syst. 32, 285–307.
Polhill, J.G., Jarvis, A., Gimona, A., Gotts, N.M. 2010a. Towards adaptive control of landscape biodiversity. International Congress on Environmental Modelling and Software (iEMSs 2010), Ottawa, Canada, 5–8 July 2010.
Polhill, J.G., Parker, D., Brown, D., Grimm, V., 2008. Using the ODD protocol for describing three agent-based social simulation models of land use change. J. Artif. Soc. Soc. Simul. 11 (2), 3. http://jasss.soc.surrey.ac.uk/11/2/3.html.
Polhill, J.G., Sutherland, L.-A., Gotts, N.M., 2010b. Using qualitative evidence to enhance an agent-based modelling systems for studying land use change. J. Artif. Soc. Soc. Simul. 13 (2), 10. http://jasss.soc.surrey.ac.uk/13/2/10.html.
Rosin, C., Campbell, H., 2009. Organification: n. The tendency to assume social and environmental orientations associated with organic production. Paper presented at the 23rd European Society for Rural Sociology Congress, Vaasa, Finland, 17–21 August 2009. Abstracts Available at: http://orgprints.org/17124/1/ESRS09 WG26.pdf.
Samhouri, J.F., Levin, P.S., Ainsworth, C.H., 2010. Identifying thresholds for ecosystem-based management. PLoS One 5 (1), e8907.
Selman, P., 2006. Planning at the Landscape Scale. In: The RTPI Library Series. Routledge, London and New York.
Simon, H.A., 1955. A behavioral model of rational choice. Q. J. Econ. 69, 99–118.
Squazzoni, F., Boero, R., 2002. Economic performance, inter-firm relations and 'local institutional engineering' in an agent-based computational prototype of industrial districts. J. Artif. Soc. Soc. Simul. 5 (1), 1. http://jasss.soc.surrey.ac.uk/5/1/1.html.
Stoate, C., Baldi, A., Beja, Boatman, N.-D., Herzon, I., van Doorn, A., de Snoo, G.R., Rakosy, L., Ramwell, C., 2009. Ecological impact of early 21st century agricultural change in Europe: a review. J. Env. Manag. 91, 22–46.
Sutherland, L.-A., 2010. Environmental grants and regulations in strategic farm business decision-making: a case study of attitudinal behaviour in Scotland. Land Use Policy 27 (2), 415–423.
Sutherland, L.-A., 2011. "Effectively organic": environmental gains on conventional farms through the market? Land Use Policy 28 (4), 815–824.
Sutherland, L.-A., Gabriel, D., Hathaway-Jenkins, L., Pascual, U., Schmutz, U., Rigby, D., Godwin, R., Sait, S.M., Sakrabani, R., Kunin, W.E., Benton, T.G., Stagl, S., 2012. The 'neighbourhood effect': a multidisciplinary assessment of the case for farmer co-ordination in agri-environmental programmes. Land Use Policy 29 (3), 502–512.
Sznajd-Weron, K., Sznajd, J., 2000. Opinion evolution in a closed community. Int. J. Mod. Phys. C 11 (6), 1157–1165.
Tilman, D., Cassman, K.G., Matson, P.A., Naylor, R., Polasky, S., 2002. Agricultural sustainability and intensive production practices. Nature 418, 671–677.
Tilman, D., Downing, J.A., 1994. Biodiversity and stability in grasslands. Nature 367, 363–365.
Vickery, J., Bradbury, R., Henderson, I., Eaton, M., Grice, F., 2004. The role of agri-environment schemes and farm management practices in reversing the decline of farmland birds in England. Biol. Conserv. 119, 19–39.
Vickery, J., Carter, N., Fuller, R.J., 2002. The potential value of managed cereal field margins as foraging habitats for farmland birds in the UK. Agric. Ecosyst. Env. 89, 41–52.
Voinov, A., Bousquet, F., 2010. Modelling with stakeholders. Environ. Model. Softw. 25, 1268–1281.
Wallis de Vries, M.F., Bakker, J.P., van der Wieren, S.E., 1998. Grazing and Conservation Management. Academic Publishers, Dordrecht, The Netherlands.
Weisskopf, W.A., 1955. The Psychology of Economics. University of Chicago Press.
Whitten, S., Gorddard, R., Knight, A., Reeson, A., Stevens, D., 2008. Designing and Testing an Outcome Focused Conservation Auction: Evidence from a Field Trial Targeting Ground Nesting Birds. 10th International BIOECON Conference on the Effectiveness and Efficiency of Biodiversity Conservation Instruments, Sidney Sussex College, Cambridge, September 28–30, 2008. http://www.ucl.ac.uk/bioecon/10th_2008/13.Whitten.pdf.
Wood, S.N., 2001. mgcv: GAMs and generalized ridge regression for R. R News 1 (2), 20–25.
Wood, S.N., 2006. Generalized Additive Models: an Introduction with R. Chapman and Hall/CRC.