International Environmental Modelling and Software Society (iEMSs)
2010 International Congress on Environmental Modelling and Software
Modelling for Environment's Sake, Fifth Biennial Meeting, Ottawa, Canada
David A. Swayne, Wanhong Yang, A. A. Voinov, A. Rizzoli, T. Filatova (Eds.)
http://www.iemss.org/iemss2010/index.php?n=Main.Proceedings

Towards Adaptive Control of Landscape Biodiversity

J. Gary Polhill[a], **Andrew Jarvis**[b], **Alessandro Gimona**[a], **Nicholas M. Gotts**[a]

[a] *Macaulay Land Use Research Institute, Aberdeen. AB15 8QH. UK.*
{g.polhill, a.gimona, n.gotts}@macaulay.ac.uk

[b] *Lancaster Environment Centre, Lancaster University. LA1 4YQ. UK.*
a.jarvis@lancaster.ac.uk

Abstract: This paper reports work using a coupled agent-based model of land use change and species metacommunity model (FEARLUS-SPOMM) to explore biodiversity incentive mechanisms using adaptive control algorithms. Existing work at the Lancaster Environment Centre has shown that such algorithms can help in analysing the adaptive capabilities of climate-related policy interventions under uncertainty. In particular, feedback control is a practical approach that mimics many recognisable features of adaptive decision making processes, including the sequential review of models and policies in the light of new observations. In a companion paper, we show that FEARLUS-SPOMM shows nonlinear dynamics in the relationship between agricultural incentives and biodiversity outcome. In this pilot study, we show that the sequential review implicit in feedback control allows information about unforeseen events to be assimilated such that biodiversity can be maintained at a desired level by adaptively adjusting agricultural incentives. We also find that Simon's 'satisficing' has its limitations in the representation of human decision making.

Keywords*: Agri-environmental incentives; model predictive control; agent-based models of land use/cover change; metacommunity models; biodiversity*.

1. INTRODUCTION

Biodiversity is conserved both for its intrinsic value and for utilitarian reasons, and many countries have committed to actions on biodiversity by signing the UN Convention on Biological Diversity[1] and enabling legislation. As land managers cannot easily benefit from the conservation of biodiversity they often need to be persuaded to adopt conservation-friendly practices through incentive schemes. However, links between incentives and outcomes are shrouded in uncertainty, because the land manager-landscape system is complex, and prey to a broad spectrum of unpredictable external influences, both natural and anthropogenic.

As with many contemporary environmental management issues, much of the uncertainty surrounding biodiversity management will be deep-seated, severely limiting the scope of management tools and practices based on predicting system behaviour. This has led to an emerging literature and practice advocating 'systems'-based approaches to managing complexity through self-adaptation [Checkland and Scholes, 1999]. However, experience has shown that self-adaptation on behalf of agents, such as farmers, would often be insufficient to engender the appropriate response; and adaptive management and adaptive interventions by policy agents are required [Holling, 1978; Secretariat of the Convention

---

1 http://www.cbd.int/

International Environmental Modelling and Software Society (iEMSs)
2010 International Congress on Environmental Modelling and Software
Modelling for Environment's Sake, Fifth Biennial Meeting, Ottawa, Canada
David A. Swayne, Wanhong Yang, A. A. Voinov, A. Rizzoli, T. Filatova (Eds.)
http://www.iemss.org/iemss2010/index.php?n=Main.Proceedings

on Biological Diversity, 2006]. Although not in vogue, hard systems methods such as offered by control systems theory may prove helpful in this area, particularly in relation to managing uncertainty [Jarvis et al. 2009].

Control theory is the mathematics of controlling dynamic systems. Here a clear distinction is made between the system to be controlled (often referred to as the 'plant') and the 'control law' facilitating this. Invariably the control law acts through feeding back signals from the plant output in order to generate its input. The mathematics determine how to manipulate the feedback characteristics in order to construct the control law so that the adjustment of the input driving variable(s) gives rise to the required behaviour(s). Control theory has been applied to decision making problems in other non-engineered systems besides that considered here, such as banking [Aoki and Nikolov, 2004] and climate change [Fiddaman, 2002]

Several adaptive mechanisms are used in control systems ranging from simple servo-mechanics to adaptive and model predictive control (see e.g. Nise [2004]). Here we report on a pilot study where a very simple integral control regime is used to mimic the adaptive capability of a Government agent when regulating biodiversity in a coupled agent-based model of land use change and species metacommunities. The Government agent's adaptive capability depends on linking adjustments in financial incentives for different land use classes to measures of landscape biodiversity through feedback. The policy action is designed to achieve a specified biodiversity target by integrating out the observed difference between biodiversity and its target value. We evaluate the performance of the adaptive policy agent and discuss wider implications for the management of complex systems under uncertainty.

2. METHOD

FEARLUS-SPOMM is a coupled agent-based model of land use change and species metacommunities. It models an environment of 25x25 land parcels, each owned by a land manager. Land managers choose land uses for the parcels they own. These affect both the money they earn from the market, and the availability of habitat for species. There are six land uses to choose from in the simulation experiments used here, of which two, labelled AL1 and GL2, provide the best habitat for a diverse assemblage of species. The metacommunity model computes the species distribution from a habitat matrix, and a Government agent monitors species distribution and issues financial incentives to achieve its biodiversity goal. Land managers use a satisficing algorithm—they do not change land uses unless their profit does not meet their aspiration. (A companion paper to this conference describes FEARLUS-SPOMM in more detail [Polhill et al., 2010].)

From the perspective of this paper, the control is used to implement the Government agent; the rest of FEARLUS-SPOMM is a black box, the input to which is a reward, u_t, and the output from which is a Shannon index of species diversity, calculated at the scale of the whole landscape H_t. The Shannon index is an entropy-based measure from information theory used in ecology. It is computed here[2] as:

$$H_t = -\sum_i \frac{o_{it}}{N_p N_s} \log \frac{o_{it}}{N_p N_s} \tag{1}$$

where i iterates over the species, o_{it} is the number of patches occupied by species i at time t, N_p is the number of patches, and N_s is the number of species. (If $o_{it} = 0$, then the summed expression is assumed to be 0, which is the limiting case as o_{it} approaches 0.) When all species occupy all patches ($o_{it} = N_p$), $H_t = \log N_s$.

---

[2] Readers with backgrounds in Ecology will be used to a slightly different formula.

International Environmental Modelling and Software Society (iEMSs)
2010 International Congress on Environmental Modelling and Software
Modelling for Environment's Sake, Fifth Biennial Meeting, Ottawa, Canada
David A. Swayne, Wanhong Yang, A. A. Voinov, A. Rizzoli, T. Filatova (Eds.)
http://www.iemss.org/iemss2010/index.php?n=Main.Proceedings

The Government agent's role is to change u_t such that H_t reaches and stays at some desired level, H_D. The simple control algorithm we use to represent this is based on the form of the relationship between H_t and u_t and is depicted schematically in figure 1. Given this relationship is unknown for FEARLUS-SPOMM we take the simplest possible approach and implement a government agent that assumes the relationship to be approximately a first order linear dynamic:

$$\hat{H}_t = a\hat{H}_{t-1} + bu_{t-1} \tag{2}$$

where $\hat{H}_t$ is the estimate for H_t. The parameter a determines the predicted relaxation time of FEARLUS-SPOMM which, based on the results from Polhill et al. [2010] was assumed to be 25 years, and hence $a = \exp(-1/25) = 0.9608$. This is equivalent to the Government agent having some prior (albeit uncertain) knowledge about the timescale of the biodiversity response to a disturbance in the system. Although it is reasonable to offer an approximate response time for FEARLUS-SPOMM, the equilibrium gain G (the value of Shannon index at $t = \infty$ for a constant unit reward) is highly non-stationary because of the richness of behaviour generated by FEARLUS-SPOMM. Therefore, the Government agent simply assumes that $G = H_E / u_0$ for any given run where E is the time the system attains a quasi-equilibrium state prior to any policy intervention. From this $b = (1 - a)G = 0.0392G$.

An appropriate feedback control structure for the first order dynamic relationship (2) is given by the following control law [see Taylor et al., 2000],

$$u_t = u_{t-1} + k(H_D - H_t) - f(H_t - H_{t-1}) \tag{3}$$

where f is a proportional control gain and k is the integral control gain (see figure 1).

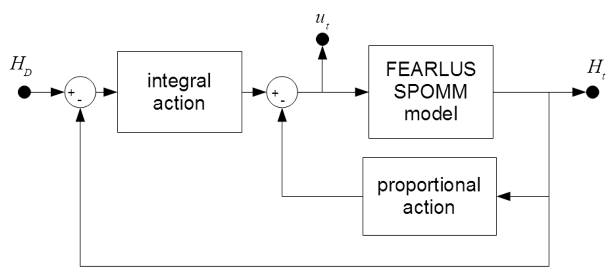

Figure 1. A schematic block diagram to show the relationship between FEARLUS-SPOMM and the control law in eq. (3) acting as the Government agent. For the specification of the gains f and k the FEARLUS-SPOMM emulation model in eq. (2) is used, whilst for the simulations FEARLUS-SPOMM itself is used.

This control law represents the Government agent's actions, generating the reward level in response to the *observed* deviation from the policy target H_D. Assimilating the observed deviations from the policy target in this way will allow the Government agent to adapt the reward in order to drive FEARLUS-SPOMM to the desired policy level in the face of uncertainty; i.e. this is the Government agent's adaptive capability. However, prior to implementation, a decision is required on the strength of the Government agent's response, as determined by k and f. These parameters are computed in relation to a and b by deciding *a priori* how fast the coupled FEARLUS-SPOMM-Government system should respond to changes in both policy targets and external disturbances. Taking the position that the Government agent wants the system to respond more quickly than the natural timescale of 25 years, we tried configuring the control law for a 10 year timescale at first, and later a 2 year timescale (see below). (Of course, these timescales cannot be expected in practice when the linear control law in eq. (3) is used to control to the nonlinear FEARLUS-SPOMM.) Since G is known to vary in time, the Government monitors the value of G prior to implementing the policy, and we nevertheless try to achieve these timescales by specifying the corresponding relationships between G, f and k through combining (2) and

(3) (assuming $\hat{H} = H$). Using a 'pole assignment' procedure [see Nise, 2004], for the 10 year case this yields $f = 3.6247 / G$ and $k = 0.2311 / G$.

The simulation is initialised using a 40 year run-in period for the land manager agents to learn. For this period $u_t = 1$. The first 30 years of the run-in period included a 'repopulation' period, in which species were reintroduced to the landscape on unoccupied patches with appropriate habitat. This is to remove effects on species levels incurred through land-use changes during the initial learning period when land managers have relatively little experience. After 40 years, the government waits for H_t to reach quasi-equilibrium (detected when $|H_t - H_{t-1}| < h$ for T consecutive time steps, where h is a parameter), allowing G (and hence b, f and k) to be estimated. $\hat{H}_t$ is initialised on H_t at this point. This marks the beginning of the ecosystem management period when the reward is adjusted as per eq. (3) using observed values of H_t taken from FEARLUS-SPOMM and $\hat{H}_t$ from eq. (2).

Our goal was to find out whether control algorithms of this kind could be used to manage the complex behaviour of FEARLUS-SPOMM, so that a desired Shannon could be reached. With 10 species in the experiment set-up, the maximum Shannon is ~2.3, but since habitat limitations of the species prevent all species being on all patches, this maximum cannot be attained. We chose two targets for the Shannon, 1.0 and 1.4, representing a modest and an ambitious target for biodiversity based on earlier model runs. After an initial exploration, we conducted two experiments, described below.

2.1 Experiment A

Table 1 shows the set up for Experiment A. There were sixteen settings exploring eight settings for the profit aspiration for the less and more ambitious Shannon index targets, each setting having 20 repetitions.

Table 1. Simulation set-ups for Experiment A. Note that some parameters (i.e H_D and profit aspiration) have multiple values. $U(1, 9)$ means the profit aspiration of land managers is assigned heterogeneously from a uniform distribution.

Parameter	Value(s)
Reward method	Reward a land manager u_t for using AL1 or GL2 on a parcel, and u_t for each neighbouring parcel with the same land use.
T	10
h	0.05
H_D	1, 1.4
u_0	1
Pole assignment	$f = 3.6247 / G$; $k = 0.2311 / G$
Proportional action	$f(\hat{H}_t - \hat{H}_{t-1})$
Run-in period	40
Market	Constant prices for each land use ('flat')
Repopulation period	30
Break even threshold	25
Profit aspiration	1, 3, 5, 7, 8, 8.5, 9, $U(1, 9)$
Decision algorithm	Make no change if aspiration met, else use case-based reasoning to review land uses on all parcels owned.
Case storage time limit	75

2.2 Experiment B

For Experiment B, we changed the decision algorithm of land managers to allow a small probability of managers reviewing land uses even if their aspirations are met, and shortened their case memory (i.e. the length of time that cases are stored in the case base). Both were aimed at making land managers more responsive to incentive change. A further change was made to the pole assignment assumptions to make the government more responsive, aiming for a 2 year response timescale. We switched to a variable market so the government was not the only driver of change in the system. The parameters in Table 2 detail the setup for two runs, each of which had 20 replications.

Table 2. Simulation set-ups for Experiment B.

Parameter	Value(s)
Reward method	Reward a land manager u_t for using AL1 or GL2 on a parcel, and u_t for each neighbouring parcel with the same land use.
T	10
h	0.05
H_D	1, 1.4
u_0	1
Pole assignment	$f = 15.122 / G$; $k = 3.949 / G$
Proportional action	$f(H_t - H_{t-1})$
Run-in period	40
Market	Variable prices for each land use ('var2')
Repopulation period	30
Break even threshold	25
Profit aspiration	1
Decision algorithm	If aspiration not met, use case-based reasoning to review land uses on all parcels owned. Else with probability 0.05, review land uses on all parcels owned, otherwise make no change.
Case storage time limit	25

3. RESULTS

3.1 Experiment A

Figure 2 shows the results of experiment A when the aspiration threshold is one, which provides the closest approximation to the emulation model in eq. (2). In each case the top graph shows a land use trajectory for the 200 time step run. One colour is used for each land use, and one line for each replication of the setting (the shade of the colours is changed between runs). The second graph shows the reward (u_t) trajectories over time, with one shade of grey for each replication and a solid line showing the *emulated* reward (i.e. the reward if everything in FEARLUS-SPOMM worked as control theory predicted). A red vertical line indicates the onset time of the policy—one line is drawn per run. The third graph shows the model-estimated and actual Shannon ($\hat{H}_t$ and H_t; pink and yellow lines respectively) for each replication, with a solid line for the *emulated* Shannon and a dashed line for the target. The bottom graph shows the species occupancy rates, one colour per species, shaded according to replication.

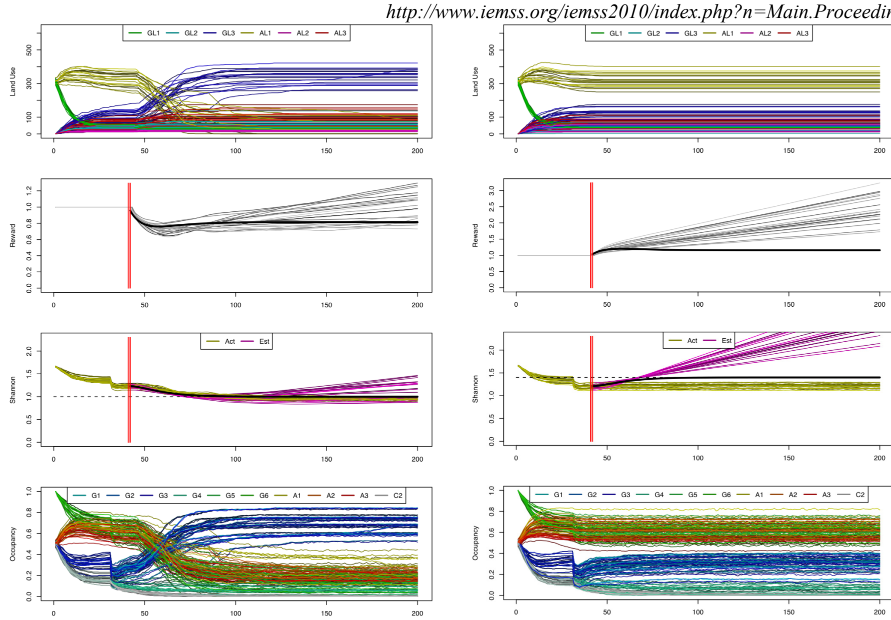

(*a*) (*b*)

Figure 2. Results from Experiment A, aspiration = 1. (*a*) $H_D = 1$; (*b*) $H_D = 1.4$. The graphs are explained in the main text.

During initialisation, whilst the farmers are learning, land use AL1 (initial reward 1.0) provides a satisficing option, and subsidises some use of GL1. At the end of the repopulation period, the persistence of C causes an immediate drop in species G1, G2 and G3. For the case where $H_D = 1.0$, the equilibrium Shannon prior to policy implementation is above target: this causes a drop in reward, turning land managers away from AL1 to GL3, and to some extent, AL3. This drops biodiversity towards the target, which is attained.

When $H_D = 1.4$, the equilibrium Shannon prior to policy implementation is below target, and the government increases the reward to encourage land managers to improve biodiversity. However, though the increase in reward enables persistence of the equilibrium Shannon land use distribution, preventing biodiversity loss, satisficing prevents land managers experimenting with new land uses that will improve biodiversity, and the target is never achieved, despite integral wind-up in the reward level.

3.2 Experiment B

Building on Experiment A, Experiment B explored the use of more responsive land managers, with a case memory of 25 rather than 75 years, and 0.05 probability of experimenting even when satisficed. This proved a double-edged sword, as Figure 3 shows. Experimentation means the more profitable but less biodiversity-friendly GL3 is adopted early in the learning period, and the equilibrium Shannon prior to policy implementation is lower than in Experiment A. However, the more responsive government agent rapidly increases reward, and once land managers forget any earlier non-satisficing experience of AL1 and GL2, these land uses get adopted. The variable market also stops managers sticking to GL3, as it will not always satisfice. The model shows sensitivity to the point at which quasi-equilibrium in Shannon index is reached and the policy adaptation started (red lines—one per run); the earliest cases all achieve and even overshoot the target, leading to a gradual decline in reward. In most of the later cases, species are lost before the government takes action, and it then vainly increases the reward to ever higher levels.

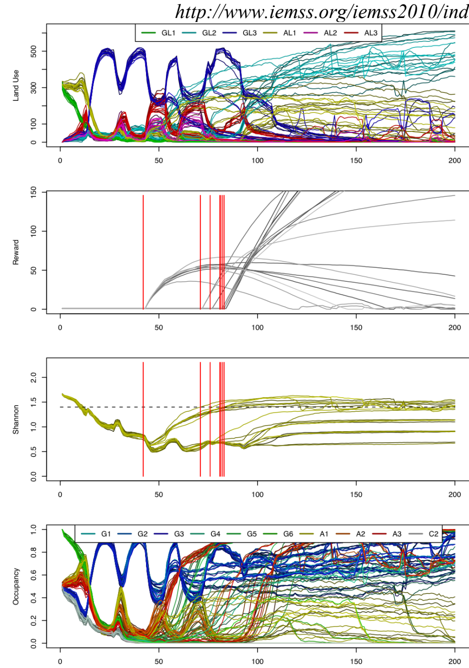

Figure 3. Results from Experiment B, aspiration = 1, H_D = 1.4.
The graphs are explained in section 3.1.

4. DISCUSSION

The results show that, under certain conditions, an adaptive policy regime based on assimilating the observed state of a system can achieve and maintain a desired level of biodiversity in FEARLUS-SPOMM. As expected, the more modest target proved easier to achieve. Although the more ambitious target was achieved only once land managers were made more responsive (Figure 3), in Experiment A (Figure 2(*b*)), it did lead to better biodiversity, and greater levels of occupancy for the more vulnerable species. Those runs where the Government agent control algorithm failed to achieve its target highlight the problems in applying a linear control framework to a nonlinear application such as this. Future work will explore including at least some consideration of the nonlinearities in the design of the Government agent's behaviour. For example, in the current runs G is constant once initialised, whereas clearly some form of online recursive estimation of G in the evolving system could and should be used to modify the control response accordingly, i.e. an *adaptive control* regime [Chotai et al., 1991].

That said, through making the Government agent significantly more adaptive (by moving the closed loop poles nearer to zero) the assimilation of the annual observed H_t combined with rapid and significant feedback response to this signal contributes significantly to the ability of the Government agent to manage system complexity. If the same is true in practice, overcoming intrinsic uncertainty when managing landscape biodiversity requires extensive monitoring of the policy relevant metric(s), short review cycles, and policy instruments that correctly identify variables land managers actually respond to.

There is an analogy between the model described and the adaptive management paradigm [Holling, 1978; Walters, 1986], but there are also differences. The analogy is in the adaptive behaviour of both land managers and government agent. An important difference is that, in the model, when rewarding for activity, the land management options needed to achieve the government goal are known with certainty. This is not necessarily the case in reality, and it is what adaptive management is trying to address. Rewarding by outcome would simulate adaptive management better. In the model, land managers experiment with different land uses, remember the outcomes, can practice a form of social learning by providing advice to their neighbours, and adapt. Rewarding by outcome, therefore, approximately simulates an adaptive management situation, although a very optimistic one, where monitoring and policy response are tightly and quickly coupled. Future work could explore rewarding by outcome rather than activity.

A note of caution from the results of Experiment B is that failure to act soon enough renders the target unachievable through loss of species. Long-term species diversity can only be maintained at the landscape scale, as colonisations from the surrounding landscape rescue local population from decline or repopulate land parcels where extinction had occurred. This work shows that a relatively high level of expenditure, of the same order of magnitude as farmers' expected profits, is needed to maintain populations in a landscape where many land managers would tend to cause their extinction through adopting profitable land uses that offer no habitat to the species of conservation interest.

Although FEARLUS-SPOMM, as currently configured, is a stylised representation of the real-world situation, these results suggest that the problem of managing complex systems for specific goals is at least sometimes tractable, contrary to a claim that goes back at least to von Mises [1922], and is based on the idea that one needs a complete model of a system to manage it intelligently. Among complex adaptive systems (CASs), with their characteristic emergence of global patterns from local interactions, path-dependence and phase shifts, we need to distinguish a subset of *managed* CASs (MCASs), where emergence interacts with monitoring and planning by agents with an overview, and extreme events prompt systemic restructuring; and within those again, *contested* MCASs (CMCASs), where multiple agents compete and form coalitions to push the system in their preferred direction.

Research is needed to discover whether the technique explored in this study can be applied when there are multiple, potentially conflicting goals for our environment, as is the case in most human-natural systems. These experiments have also assumed perfect knowledge (of the Shannon) on the part of the Government agent, and future work could investigate effects of sampling errors. A further important issue is whether control methods such as these could form the basis of an acceptable policy. Budgets for agri-environment schemes are limited, and levels of compensation to farmers tend to reflect actual and opportunity costs for activities, rather than being a 'reward'.

The results from both experiments also show a sensitivity to configuration of the land managers. Many neoclassical economic models of land manager decision making assume full fiscal rationality, an assumption that would make farmers much more responsive to policy signals than observed in these simulations. Qualitative research evidence suggests that though profit is a significant motivating factor for farmers, they do not seek to optimise economic returns—as such, Simon's [1955] 'satisficing' model may seem more appropriate, and is widely used in agent-based models of land use/cover change [Parker et al., 2008]. Whilst there are lessons here for those predicating policy decisions on utility maximisation, the results also show a weakness in the satisficing model—at least as implemented here, assuming that when aspirations are achieved managers make no changes. In the real-world, however happy farmers were with their existing income, it seems unlikely they would not respond to a large enough government incentive. Although we addressed the issue with a random probability of breaking the satisficing rule, there may be better approaches to modelling this. Possibilities include a 'temptation' threshold (where land managers whose aspirations were achieved would still look at possible alternatives, but not implement them unless the expected outcome exceeded the temptation), and relative aspiration (where land managers could be dissatisfied if neighbours are making more than they are).

5. CONCLUSION

This initial work exploring the use of control algorithms to control a nonlinear model has demonstrated that even a simple linear control algorithm can have at least partial success in achieving its target. We anticipate that future work implementing more sophisticated control algorithms will prove more successful, and in a wider range of scenarios. Insofar as FEARLUS-SPOMM exhibits complex dynamics, our results suggest that complex adaptive systems can be governed; albeit that this may require constant monitoring, and swift and decisive adaptation of incentives to maintain targets. Whilst some of the literature on complex systems has tended to highlight self-organisation and emergence—suggesting 'light-touch' government in contrast to 'command-and-control', our results also show that timely government intervention is necessary to protect species from local extinction.

ACKNOWLEDGMENTS

This work was funded by the Scottish Government Rural and Environmental Research and Analysis Directorate. Contributions to FEARLUS-SPOMM software 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.

REFERENCES

Aoki, K. and K. Nikolov, Rule-based monetary policy under central bank learning, *Working paper no. 235*, *Bank of England*, 2004.

Checkland, P. and J. Scholes, *Soft Systems Methodology in Action,* Wiley, Chichester, 1999.

Chotai, A., P.C. Young, M.A. Behzadi, Self-adaptive design of a non-linear temperature control system, special issue on self-tuning control. *IEE Proceedings Part D*, *38*, 41–49, 1991.

Fiddaman, T.S., Exploring policy options with a behavioral climate-economy model, *System Dynamics Review*, *18*, 243–267, 2002.

Holling, C.S. (ed.), *Adaptive Environmental Assessment and Management*, Wiley, Chichester, 1978.

Jarvis, A., D. Leedal, C.J. Taylor, P.C. Young, Stabilizing global mean surface temperature: A feedback control perspective. *Environmental Modelling and Software*, *24*, 665–674, 2009.

Nise, N.S., *Control Systems Engineering*, John Wiley, New York, 2004.

Parker, D.C., A. Hessl, and S.C. Davis, Complexity, land-use modelling, and the human dimension: Fundamental challenges for mapping unknown outcome spaces. *Geoforum*, *39*(2), 789–804, 2008.

Polhill, J.G., A. Gimona, and N.M. Gotts, Analysis of incentive schemes for biodiversity using a coupled agent-based model of land use change and species metacommunity model, *International Congress on Environmental Modelling and Software (iEMSs 2010), Ottawa, Canada, 5–8 July 2010*.

Secretariat of the Convention on Biological Diversity, *Global Biodiversity Outlook 2.* Montreal, 2006.

Simon, H.A., A behavioral model of rational choice. *The Quarterly Journal of Economics LXIX*(February), 99–118, 1955.

Taylor, C.J., A. Chotai, P.C. Young, State space control system design based on non-minimal state variable feedback: further generalization and unification results. *International Journal of Control*, *73*, 1329–1345, 2000.

von Mises, L., *Die Gemeinwirtschaft: Untersuchungen über den Sozialismus,* Gustav Fischer Verlag, Jena, 1922, 1932, English translation published as *Socialism: An Economic and Sociological Analysis,* Jonathan Cape, London, 1936.

Walters, C., *Adaptive Management of Renewable Resources.* Macmillan, New York, 1986.