Model-based optimization of agricultural profitability and nutrient management: a practical approach for dealing with issues of scale

Model-based optimization of agricultural profitability and nutrient management: a practical approach for dealing with issues of scale

Alice E. Milne, Kevin Coleman, Lindsay C. Todman & Andrew P. Whitmore

Addresses

Alice E. Milne, Kevin Coleman, & Andrew P. Whitmore

Sustainable Agriculture Sciences, Rothamsted Research, Harpenden, Herts, AL5 2JQ, UK

Lindsay C. Todman

School of Agriculture, Policy and Development, University of Reading, Berks, RG6 6AR, UK

Corresponding author Alice E Milne (email: alice.milne@rothamsted.ac.uk , Tel: +44 1582 938 380)

ORCID ID's

Alice Milne (0000-0002-4509-0578)

Kevin Coleman (0000-0002-9640-1479)

Lindsay Todman (0000-0003-1232-294X)

Andrew Whitmore (0000-0001-8984-1436)

Acknowledgements

Rothamsted Research receives grant aided support from the Biotechnology and Biological Sciences Research Council (BBSRC) of the United Kingdom. This research was funded the Biotechnology and Biological Sciences Research Council (BBSRC) Institute Strategic Programme (ISP) grants, "Soils to Nutrition" (S2N) grant number BBS/E/C/000I0330, the joint Natural Environment Research Council (NERC) and the Biotechnology and Biological Sciences Research Council (BBSRC) ISP grant "Achieving Sustainable Agricultural Systems" (ASSIST) grant number BBS/E/C/000I0130 ASSIST, using facilities funded by the BBSRC, and by a DEFRA and EU collaborative project "Targets for Sustainable And Resilient Agriculture" (TSARA), received as part of the FACCE-JPI Surplus initiative.

Abstract To manage agricultural landscapes more sustainably we must understand and quantify the synergies and trade-offs between environmental impact, production and other ecosystem services. Models play an important role in this type of analysis as generally it is infeasible to test multiple scenarios by experiment. These models can be linked with algorithms that optimise for multiple objectives by searching a space of allowable management interventions (the control variables). Optimisation of landscapes for multiple objectives can be computationally challenging, however, particularly if the scale of management is typically smaller (e.g. field-scale) than the scale at which the objective is quantified (landscape scale) resulting in a large number of control variables whose impacts do not necessarily scale linearly. In this paper, we explore some practical solutions to this problem through a case study. In our case study we link a relatively detailed, agricultural landscape model with a multiple-objective optimisation algorithm to determine solutions that both maximise on profitability and minimise greenhouse gas emissions in response to management. The optimisation algorithm combines a non-dominated sorting routine with differential evolution, whereby a "population" of 100 solutions evolve over time to a Pareto optimal front. We show the advantages of using a hierarchical approach to the optimisation, whereby it is applied to finer scale units first (i.e. fields), and then the solutions from each optimisation are combined in a second step to produce landscape-scale outcomes. We show that if there is no interaction between units then the solution derived using such an approach will be the same as the one obtained if the landscape is optimised in one step. However, if there is spatial interaction, or if there are constraints on the allowable sets of solutions then outcomes can be quite different. In these cases, other approaches to increase the efficiency of the optimisation may be more appropriate – such as initialising the control variables for half of the population of solutions with values expected to be near optimal. Our analysis shows the importance of aligning a policy or management recommendation with the appropriate scale.

Keywords Landscape modelling, trade-offs, synergies, environmental impact, multiple-objective optimisation

Introduction

Agricultural landscapes provide our food, contribute to the way natural resources are managed, and provide areas for recreation and public wellbeing (Westmacott and Worthington 2006). Pressures to increase food production have led to many unsustainable agricultural practices which can degrade the soil, reduce water quality, increase the likelihood of flooding, impact biodiversity and result in the emissions of greenhouse gases (Bennett et al. 2009; Seppelt et al. 2016; Tilman et al. 2002). Mitigating anthropogenic impacts on the environment and global food security are hence two major challenges, and identifying and exploiting synergies between these should result in social, economic and ecological benefits (Cramer et al. 2017). Sound landscape management strategies are therefore essential for the long-term sustainability of agriculture, and so it is not surprising that there is an increasing amount of research into how we should manage agricultural landscapes to fulfil multiple objectives aligning to production and environmental quality (Kennedy et al. 2016; Groot et al. 2018; Verhagen et al. 2018; Fischer et al. 2017; O'Farrell and Anderson 2010). This ambition, however, inevitably involves trade-offs between conflicting objectives (Howe et al. 2014).

In much of the research done on landscape design and management, a recurring theme is the need to understand and quantify the synergies and trade-offs between environmental impact, production and other ecosystem services (Gourevitch et al. 2016; Howe et al. 2014; Kennedy et al. 2016). Approaches that rely on data and measurement are hampered by the fact that it is often infeasible to experiment at the scales (both spatial and temporal) appropriate to how best to manage landscapes. Not surprisingly therefore, computer simulation models have an important role to play in filling the large gaps between what we need to know and what is available from measurements. Many approaches rely on scenario analysis whereby various management strategies or policies are tested through simulation. A second approach, which we explore here, is to link a model that describes the impact of management on an agricultural landscape with an optimisation algorithm, and so determine the sets of inputs to the model (known as "the control variables") that maximise the desired outcomes in the model. These outcomes are framed as an "objective function" and could be any combination of profit and measures of environmental impact. The optimisation algorithm seeks to maximise (or minimise) the objective function by efficiently searching the allowable ranges of the control variables. Linking models of ecosystems services with optimisation algorithms to elucidate mechanisms to fulfil multiple objectives is becoming increasingly popular. Kennedy et al. (2016) used models of agricultural profit, biodiversity and freshwater quality linked to an optimisation algorithm to investigate trade-offs under various land use scenarios. Their objective function was formed from a weighted sum of the individual objectives. They demonstrate the advantages of considering multiple objectives when optimising landscape management strategies, over optimisation based on production or profit alone. Their analysis showed that through joint planning for economic and environmental goals at a landscape-scale, Brazil's agricultural sector could expand production and still meet regulatory requirements, while maintaining biodiversity and ecosystem service provision. Others have advocated the use of multi-objective optimisation, whereby the optimisation algorithm is used to determine Pareto optimal fronts of multiple objectives. The Pareto front describes the trade-off between objective variables such as yield and biodiversity where it is not possible to improve outcomes for one variable without impacting another adversely. For example, Verhagen et al. (2018) present a multi-objective optimization of on- and off-farm agri-environment measures to maximise fruit production, potential habitats for endangered species, and landscape aesthetics whilst minimising loss of pasture production. The models that they use include lookup tables as well as more complex approaches. Groot et al. (2018) present a landscape modelling framework for multi-scale spatially explicit analysis of trade-offs and synergies among ecosystem services. They include a multiple objective optimisation to determine trade-offs between ecosystems services that may be estimated from simple relationships or more complex models. Teillard et al. (2017) apply multiple objective optimization to determine how the spatial planning of agricultural intensity allocation could improve on both food production and the diversity of farmland birds on a national scale. Their optimisation considers the whole of France with control variables applied at the scale of small agricultural regions(590 regions were used in the analysis).

A challenge that frequently arises in model-based optimisation of landscape management relates to the scale at which the control variables should be applied. Typically, decisions on landscape management are taken at relatively fine scale: field scale or finer, whereas the objectives we wish to optimise are at the scale of the landscape. This discrepancy in scale can lead to an intractably large number of control variables. For example, if we consider the management of fertilizer on a field-by-field basis across a landscape, (even without consideration of any other control variable), the number of fertiliser controls can be or the order of hundreds to thousands. This number of control variables makes convergence to an optimal solution unlikely. In this study, we have explored some practical solutions to such problems where there is a discrepancy between the scale of implementation of a control and the scale of a desired outcome. To explore these, we linked a relatively detailed model that describes an agricultural landscape (Coleman et al. 2017) with a multiple objective optimisation algorithm. The example that we consider is how to manage a landscape for improved nutrient use efficiency (i.e. reducing nutrient losses through greenhouse gases and leaching whilst maintaining good productivity across the landscape). Here we consider the implications of taking a hierarchical approach to this type of problem, whereby we optimise the management decisions made on a field-by-field basis first, and then combine these in subsequent steps. We explore the conditions under which such an approach would be beneficial, and where it would not. We work with a simulated landscape based on a 1km x 1km square of arable land in the UK, and demonstrate that our approach can provide solutions to this large-scale problem. In particular, we explore the implications such an approach has when our landscape has substantial spatial interaction or when there are conditions (or constraints) on the allowable set of solutions. We also consider approaches that may improve on the rate of convergence of our optimisation. We conclude with some broad recommendations and discuss how more complex scenarios could be approached.

Method

We linked the Rothamsted Landscape model (RLM) with a multiple objective optimisation algorithm to explore practical approaches to scaling up model-based optimisation of landscape management. We start by describing the model and case studies before going on to describe the optimisation algorithm and the strategies we investigated to more efficiently explore the search space for optimal management.

Landscape model

The Rothamsted Landscape model, RLM, (Coleman et al. 2017) simulates the effect of fertilizer management on profit (calculated as the difference between income from yield and the costs associated with fertilizer and its application), yield and the environment. This model operates at a daily time step and simulates the essential processes of soil, water, crop growth and biodiversity for agricultural landscapes in the UK (Fig. 1). The crop model is a generic plant growth model based on LINTUL (Wolf 2012; Shibu et al. 2010). The model has been parameterised for 20 crops including major cereal crops, grass, potatoes sugar beet, and onions. The RLM also has an arable weed component that simulates 136 weed species (Metcalfe et al. 2019).

The simulation of soil-water dynamics uses a capacity based approach (Addiscott and Whitmore 1991) where the capacity of each layer depends on soil texture, soil organic matter and bulk density. Water is available for crop uptake and is lost through percolation, runoff, evaporation and transpiration. The soil organic carbon (SOC) dynamics are based on the Rothamsted carbon model, RothC, (Coleman and Jenkinson 2014) Soil organic nitrogen (SON) and soil organic phosphorus (SOP) are modelled in a similar way to the SOC dynamics, both SON and SOP have the same pool structure as the active SOC pools. Soil mineral nitrogen comprises ammonium ($\mathrm{NH_4^+}$) and nitrate ($\mathrm{NO_3^-}$) and is input through atmospheric deposition, and inorganic fertilizer application as well as mineralisation from soil organic matter. When organic amendments are added, N enters the soil inorganic nitrogen pools by mineralisation. Mineral nitrogen may be taken up by the crop and is lost through runoff , leaching ($\mathrm{NO_3^-}$ only) and emissions from the soil. Mineral phosphorus is added as fertilizer may be taken up by the crop and can be lost through runoff. Full details are given in Coleman et al. (2017).

The RLM is spatially explicit. This is achieved by considering the area to be modelled as a grid of cells where each cell represents a field or part of a field (depending on the scale of interest). To initialise the landscape, soil properties are set in each cell and the soil water content of each cell is set to field capacity. Within each cell, we model crop growth, the dynamics of soil water, SOC, SON, SOP, changes in bulk density and nutrient (i.e. inorganic N and P) flows on a daily time step. Water and nutrients can move laterally between cells as runoff, as well as vertically though the soil profile, as drainage. The landscape model is modular with the main infrastructure (calls to subroutines and data handling) written in C++ and other modules (crop growth, soil and water processes, weed dynamics and livestock) written in either Fortran or C++.

Modelled landscape scenarios

To explore the basic principles of scaling up the optimisation to landscape scale, we considered three different scenarios. Each scenario was run over nine seasons. First, we considered a simple 1x2 grid with no spatial interaction and a crop of continuous winter wheat in both cells. We assumed each grid-cell to be of size 100m x 100m (which equates to 1ha). The soil properties for each cell were based on the soil found in two fields in Silsoe, Bedfordshire, UK, which we examined in a previous study (Lark et al. 2004). We chose these soils because they are contrasting yet found close to one another, making our simplistic scenario plausible yet diverse enough for optimal solutions to vary between cells. The soil conditions for the two fields are shown in Table 1. The model requires initial conditions for soil properties three layers deep, but we only had measurements for the top layer (Table 1). We based the soil conditions for the other two layers on some broad assumptions. We assumed that the sand, silt, clay, and pH took the same value throughout the three layers. We assumed that the organic carbon in the second (23–46cm) and third (46–69cm) was 50% and 25% of the value for the top layer respectively. The bulk density for layers 2 and 3 was estimated using the (Rawls 1983) nomogram which uses values of texture and organic carbon to estimate bulk density. As our aim was to simulate plausible field conditions, and not specifically evaluate the two fields from Silsoe, we considered these assumptions acceptable. The second scenario was identical except that this time we assume that there is a 5% slope and that water and nutrients flow laterally from “Field 2” to “Field 1” from where it runs off and is accounted for in the drainage water.

In our third scenario, we consider a more realistic landscape using a larger 10x10 grid (cell size 100m x 100m) which is based on a 1 km x 1 km area of the UK in cereal production. For this scenario, we assume that each field is in a three-year or six-year rotation somewhat typical of a rotation found in the UK (wheat–beans–wheat–barley–wheat–oilseed rape or wheat–wheat–oilseed rape). The point in the rotation that each field is started with varies across the landscape (see Fig. 2). Although we had information on the topography of this area of the UK, we did not have detailed information on soil type. We therefore assumed that the soil properties had a similar range to those we used in our 1x2 grid and allowed the properties to vary in relation to elevation with lighter sander soils associated with higher cells and heavier soils associated with lower points.

The Optimisation Algorithm

We coupled the simulation model with an optimisation algorithm to determine Pareto optimal fronts between multiple objectives defined in terms of outputs from the model. For each management unit (e.g. field), the control variables comprised the amount of inorganic N-fertilizer applied, the amount of inorganic P- fertilizer applied and the amount of organic amendment, farmyard manure (FYM), applied. Because these are control variables, we do not fix the amounts of fertilizer a- priori as one would in scenario analysis, rather we let the optimisation algorithm search the allowable space for the amounts that optimise the objective function. In the optimisation, fertilizer-N can be applied on any of nine dates starting from the sowing date or the 14th February (whichever is later) and then every ten days after. This is a pragmatic way to include variable timing in the optimisation, without explicitly adding timing as an additional control variable (Parsons and Beest 2004), as we expect that many of the nine application rates will be zero. The timings of fertilizer-P and FYM are fixed to a week before sowing and the sowing date, respectively. The N fertilizer variables were bounded between 0 and 300 kg N ha^{-1} per application, P fertilizer between 0 and 100 kg P ha^{-1}, and the FYM between 0 and 3 t C ha^{-1}. So that our results are straightforward to interpret, we restrict the number of objectives to two: profit (£ ha^{-1}) and nitrous oxide emissions (expressed in kg CO_2-equivalent ha^{-1} year $^{-1}$ where we assume a conversion factor of 298 CO_2 to N_2O).

The profit function is calculated as sum of the yield multiplied by the price of the crop each season, minus the total cost of applying fertilizer, which is made up of an application cost (£ per application) and the price of the N and P applied (£). This is divided by the number of seasons (9 as stated above) to give the average profit. In the simulations shown here FYM is assumed to be free but does incur an application cost.

The optimisation algorithm that we used combines a non-dominated sorting routine from NSGA-II (Deb et al. 2002) with differential evolution (Storn and Price 1997). These algorithms were coded in C++ and linked directly to the RLM code. Our aim is to use the optimisation algorithm to define a Pareto front of optimal solutions. For this we maintain a population of 200 solutions. Initially, the optimisation algorithm randomly generates values for the control variables for each member of the initial population. In our case this is 200 sets that define the rates on N, P and FYM fertilizer to be applied. These management strategies are then implemented in the model resulting in 200 sets of values for the optimisation objective function (in our case profit and greenhouse gas emissions). The non-dominated sorting identifies the options that result in the 'best' 100 objectives, i.e. those that are non-dominated in the sense that no other point has both a greater profit *and* a lower rate of emissions. A point is said to be dominated by another if it is worse for every single objective (for example, for a two-dimensional Pareto front describing the trade-off between profit and greenhouse gas emissions a scenario whereby profit was p_1 and emissions were g_1 would be dominated by another if $p_2 > p_1$ and $g_1 > g_2$ where p_2 and g_2 represent the profit and emissions from the second scenario). The differential evolution algorithm then combines aspects of the management options that led to non-dominated objectives (i.e. takes two sets of control variables and swaps some of the elements between the two), along with some randomisation to identify new management options that could potentially perform even better and forms a new population of 200 from which the best 100 are again selected. The process is iterated in directions that the differential evolution algorithm suggests will be an improvement, until the results converge and produce a similar Pareto front with each iteration.

Landscape optimisation Strategies

We compared four strategies for optimising landscape units for our 1x2 grid scenarios. In the first approach (Strategy 1), we optimised the landscape units separately and produced Pareto frontiers for each landscape unit. These frontiers were then combined in a second step to produce an optimal frontier for the landscape (Todman et al. 2019)). Any interaction between the two units was therefore neglected. In the second approach (Strategy 2), we assumed that the same fertilizer management should be applied to all landscape units and optimised accordingly (that is to say, the landscape was optimised at a larger scale). In the third approach (Strategy 3), we optimised the landscape in one step, assuming that each unit was managed separately. For this third approach we started the optimisation with a population set where the control variables were generated randomly. In our fourth approach (Strategy 4) we initialised half of the population of controls using the solutions generated when we optimised the units separately. In control theory terms, we “seeded” part of our population of controls with values likely to be near optimal. We also explored the difference between sets of solutions generated using Strategies 1 and 3 when a condition that the amount on maximum amount of N that could leach (an arbitrarily set threshold of 20 kg N ha^{-1}) was imposed on the allowable set of solutions. For each approach, we determined the number of iterations before the solution converged and the time taken for convergence. Based on our findings from this investigation, we applied the optimisation to the larger more realistic 10x10 landscape.

Results

Optimisation without condition on the maximum amount of N leached

The number of iterations for the solutions to converge and the times taken are shown in Table 2. We note that the absolute times to converge depend on the computer hardware, but the relative lengths of timings are informative. The time taken for the two single fields to converge, was less than half of that taken for the two-cell grid to converge. When the population of solutions was partially initiated with solutions from the single cell optimisations this time reduced to be similar to that taken for the single cell optimisation. However, the time to optimise the single cells should be also accounted for in this scenario.

There was no substantial difference in the time taken for the 1 x 2 grid with spatial interaction to converge compared with the time taken for the grid without spatial interaction.

The time for the case where management is assumed to be the same across the 1x2 landscape was similar, to the single cell solutions. Based on these results we optimised our 10 x 10 cell landscape using solutions from single cell optimisations to initialise half the population of solutions. The other half of the population was initialised randomly. We found that the population of solutions were able to converge to a frontier, although this took a substantial amount of time (see Table 2).

The optimised solutions for the two separate fields show distinct populations (Fig 3) that relate to various types of fertilizer treatments. In both fields, there is a population of solutions where only P fertilizer is applied (shown in green). These solutions are characterised by low profit and low emissions. In fact, in these solutions applying P fertilizer is not cost effective and only has advantage because the slight increase in yield that it causes results in more N going into the plant and so less lost as N_2O emissions. The populations shown in blue related to solutions where only fertilizer-N is applied. Increasing N fertilizer results in larger and more profitable yield, but emissions of N_2O increase. Field 1 has an additional population of solutions (shown in orange) these relate to applications of FYM. This source of fertilizer is cheaper than mineral N so gives greater profit in Field 1 but also result in greater emissions. There are no equivalent sets of solutions for Field 2. This difference is due to the soil. The soil in Field 1 has a greater content of clay and so additions of FYM have greater impact on improving the bulk density of the soil and hence water holding capacity than Field 2. The crop, therefore, suffers less water stress. The optimised solutions for the 1x2 grids are shown in Fig. 4–6. Combining the two sets of optimal solutions shown in Fig. 3 gives the set of solutions shown in Fig. 4. If there is no interaction between fields, the Pareto optimal frontier of this set of solutions is the same that is given by optimising the landscape as a whole (shown by the black discs in Fig. 4) i.e. the solution of a problem with, in this case, twice as many control variables. If, however, there is interaction between the landscape units (i.e. fields) then the two-step optimisation process does not reach the same solution as when the landscape is optimised in one stage (Fig. 5). We also optimised the landscape with the assumption that management was uniformly applied (Fig 6). Not surprisingly, improvements in both emissions and profit can be made if the control is allowed to vary at the finer scale (single cell) rather than be uniformly applied across soils that are substantially different. The improvements, however, are small for the solutions that relate to mineral nitrate application (on average £30 ha^{-1} $year^{-1}$ and 30 kg CO_2 eq ha^{-1} $year^{-1}$) compared with the solutions where FYM or P-fertilizer is applied. In particular, the two solutions with the largest emissions derive from occasions where FYM is applied in both fields.

Optimisation with constraints

When the constraint was imposed at the larger scale (i.e. when the cells were optimised together rather than separately and then the solutions merged) more solutions were viable (Fig. 7) as N leached in from one cell could be compensated for by smaller losses from the other cell. In particular, this affected the profitability that could be achieved with the given constraint.

Optimisation of 10 x 10 cell landscape

The 10 x 10 cell grid converged to a frontier with similar (but less distinct) populations of solutions to that observed for the 1x2 grid (Fig. 8). That is to say, there was a distinct set of solutions that related to P-fertilizer only, which were characterised by low emissions and small profit. A second cluster was characterised by moderate rates of N- and P-fertilizer but little to no FYM. The final set solutions comprised solutions with larger additions of all fertilizer types.

Discussion

Optimisation of landscapes for multiple objectives is complex particularly if the management controls available are applied at fine scale, for example, field scale management. In such cases, the number of control variables can become infeasibly large and it may no longer be possible to use an optimisation algorithm. We have explored some practical solutions to approach such a difficulty.

One way to reduce the number of control variables used in any single optimisation step is to take a hierarchical approach whereby the optimisation is applied to finer scale units, for example field scale, and then the solutions from each optimisation are combined in a second step. We show that if there is no interaction between units then the solution derived using such an approach will be the same as the one obtained if the landscape is optimised in one step, provided of course that neither approach gets stuck in a local minimum. A hierarchical approach could also be used if the number of control variables within each spatial unit is large. In this case the control variables could be grouped into sub-groups such that the expected interaction between the control variables within each sub-group is large and the interaction between the sub-groups of control variables is minimal. The advantages of the hierarchical approach are clear: the number of control variables used to determine the solution of a single unit is far fewer and the search space is therefore far less complex meaning that the chances of getting stuck in a local minimum are greatly reduced. Secondly the process of optimising the landscape can be parallelised reducing the time taken to reach a solution.

A second strategy is to apply the control variables at a larger scale than an individual unit. We showed that this had clear advantages in the time taken to converge to a solution and can reduce complexity enormously. To use this strategy wisely, some form of pre-clustering algorithm should be applied to the landscape to group similar landscape units together and apply the controls at the scale of these groupings.

The problem is less straightforward if there are interactions between cells. In these cases, the optimal solution discovered using the hierarchical approach is likely to come to a different solution compared with the one found when the landscape is optimised in one step. As we demonstrate, there is also an issue with the hierarchical approach if we apply conditions on the set of allowable solutions at a scale greater than the size of the unit that we optimise. In the example that we consider, we imposed a condition that N leaching could not exceed a specified limit per hectare. If this limit is imposed at the scale of the field (or unit cell) then we miss solutions that exploit the opportunity to exceed the limit in certain cells, compensating for this by imposing much lower levels than the threshold in others. This is analogous to imposing a regulation on water quality at catchment scale despite the fact pollutants are generally managed at field scale.

Where it is not possible to take a hierarchical approach to the optimisation, it may be advantageous to strategically “seed solutions". This is particularly appropriate with the genetic algorithm that we used as it is possible to pre-populate a proportion of the solutions leaving the remaining solutions random and hence maintaining the potential for a broad group of optimal solutions. In our case, we pre-seeded 50% of our controls with values that led to optimal solutions in the individual units. Because we seek to optimise multiple objectives, we needed to ensure that these composite sets were similarly sorted from objectives that favoured lower emissions to those that favoured profit so that the composite solutions were closer to the feasible frontier than one we might expect from random. This approach, admittedly has drawbacks. It is time consuming to set up the initial solution set, and such a construction is more likely to lead the algorithm to get stuck in local minima compared with truly random initial conditions. This risk, however, could be minimised by using different seeding strategies such as using a small percentage of seeded solutions, or seeded from partial solutions (e.g. with the controlss for one spatial unit, but with randomised controls for all other spatial units). Further options for this initial population could also be developed based on the ideas of stakeholders or by generating possible scenarios, as has been done elsewhere (Hu et al. 2015). Here, however, we demonstrated that a simple seeding approach can make it possible to optimise relatively large and complex landscape units.

In the case study we considered we looked at two objectives to simplify our exposition, however it is straightforward to include more. With this particular model we can include up to six. The objectives may be synergistic, whereby an improvement in one is positively correlated with another, or more interestingly there maybe trade-offs between pairs of objectives. The two that we chose to use demonstrate a trade-off between production and environment – with little obvious synergy. To increase profit we must fertilize (accepting there is some economic optimum) but this is often to the detriment of the environment. However, one interesting interaction picked up by the model was that if we increase P-fertilizer, potential yield can increase allowing more N to be taken up by the plant resulting in smaller N_2O emissions; however, the application of P was not cost effective in this case. This relationship between yield potential and fertilizer demand is widely acknowledged (Kindred et al. 2015). Hughes et al. (2011) observed that the use of crop protection chemicals reduces greenhouse gas emissions per unit N applied. The practical message for farmers is that alleviating limitations on yield potential increases nutrient use efficiency which can lead to larger yields and reduced N-losses.

Interestingly, the clustering solutions as described by Todman et al. (2019) shows that they fall into two or three different fertilization strategies (depending on soil type) that group somewhat along the trade-off curve (i.e. result in similar outcomes). This demonstrates the power of the optimisation approach, in that it elucidates clear patterns which are helpful when evaluating environmental response to management. In particular, we saw that on the clay soil additions of FYM can increase yield substantially but at the cost of increased emissions. This highlights the potential for increasing the value of the objectives by allowing for finer-scale management solutions (as illustrated by Fig. 6), and the importance of aligning management recommendation with the appropriate scale. Indeed, there is potential for model-based optimisation (such as that presented here) to aid farmers in decisions related to resource allocation to maximise nitrogen use efficiency.

We also showed that the scale at which a constraint or condition is applied can have a large impact on the sets of allowable solutions (Fig. 7). This has implications for policy as it demonstrates the importance of aligning a policy with the appropriate scale. Policy-makers might relax the requirement for water draining from each field to be of satisfactory quality if aggregate water from several fields meets standards. In practical terms, our analysis revealed that there is potential value in devising policy restrictions where cooperation is both allowed and encouraged. Indeed, model-based landscape optimisation, offers a key tool for policy to determine where cooperation and more flexible approaches to regulatory mitigation strategies could enhance the multiple objectives we seek to fulfil with landscape management.

The methodology described here can be extended to explore the implications of landscape management on wider sets of ecosystems services and natural capital: in particular provisioning (food production and fresh water), regulating services (climate, flood, pest and disease regulation, and pollination), and biodiversity. We can capture these facets as objectives in our optimisation producing a multidimensional surface on which each point represents a set of management options that are optimal in some way. Whilst, infield management of crop pests might aim to reduce some aspects of biodiversity (for example weed control), at the larger landscape scale we typically aim to enhance biodiversity. Indeed, there is a growing interest in the role of biodiversity and the services it generates such as natural crop protection as well as its role in cultural ecosystems service provision (Letourneau et al. 2009).

Crop choice (including grazing systems) and the associated concepts of in-field rotation are the key drivers of landscape outcomes (production and environmental impacts) and so offers an obvious yet complex set of control variables (Dury et al. 2012; Sethi et al. 2006). Varying crops intelligently in the landscape, including some sort of set-aside to enhance biodiversity, should work well cooperatively where a high-yielding but polluting crop is matched with poorer-yielding but cleaner companion. This concept sits uncomfortably with modern pressures and ways of working, however, such as block cropping and contract management that deliver economies of scale. In practice, rotation may remain a stratagem that continues to deliver sustainability over time. Model-based optimisation of such a problem could offer great insights but the complexity is enormous and so developing a hierarchical approach, similar to that described here would almost certainly be essential.

Conclusion

Model-based landscape optimisation is hampered by the fact that management interventions occur at a relatively fine-scale meaning that the number of control variables can become intractably large. We show that if there is limited interaction between spatial units (e.g. fields) then a hierarchical approach, whereby the optimisation is applied to finer scale units before being combined in a second step, can be used to advantage. If there are spatial interactions between units or constraints are applied at the landscape scale, then this approach may not be appropriate. Model-based landscape optimisation can reveal opportunities for more efficient management by farmers and for improvements to policy interventions aimed at mitigating the environmental impacts of landscape management.

References

Addiscott, T. M., & Whitmore, A. P. (1991). Simulation of solute leaching in soils of differing permeabilities. *Soil Use and Management, 7*(2), 94-102.

Bennett, E. M., Peterson, G. D., & Gordon, L. J. (2009). Understanding relationships among multiple ecosystem services. *Ecology Letters, 12*(12), 1394-1404, doi:10.1111/j.1461-0248.2009.01387.x.

Coleman, K., & Jenkinson, D. S. (2014). *RothC - A Model for the turnover of carbon in soil: Model description and users guide (updated June 2014)*. Harpenden, UK: Lawes Agricultural Trust.

Coleman, K., Muhammed, S. E., Milne, A. E., Todman, L. C., Dailey, A. G., Glendining, M. J., et al. (2017). The landscape model: A model for exploring trade-offs between agricultural production and the environment. *Science of the Total Environment, 609*, 1483-1499, doi:https://doi.org/10.1016/j.scitotenv.2017.07.193.

Cramer, W., Egea, E., Fischer, J., Lux, A., Salles, J. M., Settele, J., et al. (2017). Biodiversity and food security: from trade-offs to synergies. *Regional Environmental Change, 17*(5), 1257-1259, doi:10.1007/s10113-017-1147-z.

Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *Ieee Transactions on Evolutionary Computation, 6*(2), 182-197, doi:http://dx.doi.org/10.1109/4235.996017.

Dury, J., Schaller, N., Garcia, F., Reynaud, A., & Bergez, J. E. (2012). Models to support cropping plan and crop rotation decisions. A review. *Agronomy for Sustainable Development, 32*(2), 567-580, doi:10.1007/s13593-011-0037-x.

Fischer, J., Meacham, M., & Queiroz, C. (2017). A plea for multifunctional landscapes. *Frontiers in Ecology and the Environment, 15*(2), 59-59, doi:10.1002/fee.1464.

Gourevitch, J. D., Hawthorne, P. L., Keeler, B. L., Beatty, C. R., Greve, M., & Verdone, M. A. (2016). Optimizing investments in national-scale forest landscape restoration in Uganda to maximize multiple benefits. *Environmental Research Letters, 11*(11), 114027, doi:10.1088/1748-9326/11/11/114027.

Groot, J. C. J., Yalew, S. G., & Rossing, W. A. H. (2018). Exploring ecosystem services trade-offs in agricultural landscapes with a multi-objective programming approach. *Landscape and Urban Planning, 172*, 29-36, doi:10.1016/j.landurbplan.2017.12.008.

Howe, C., Suich, H., Vira, B., & Mace, G. M. (2014). Creating win-wins from trade-offs? Ecosystem services for human well-being: A meta-analysis of ecosystem service trade-offs and synergies in the real world. *Global Environmental Change-Human and Policy Dimensions, 28*, 263-275, doi:10.1016/j.gloenvcha.2014.07.005.

Hu, H. T., Fu, B. J., Lu, Y. H., & Zheng, Z. M. (2015). SAORES: a spatially explicit assessment and optimization tool for regional ecosystem services. *Landscape Ecology, 30*(3), 547-560, doi:10.1007/s10980-014-0126-8.

Hughes, D. J., West, J. S., Atkins, S. D., Gladders, P., Jeger, M. J., & Fitt, B. D. L. (2011). Effects of disease control by fungicides on greenhouse gas emissions by UK arable crop production. *Pest Management Science, 67*(9), 1082-1092, doi:10.1002/ps.2151.

Kanter, D. R., Schwoob, M. H., Baethgen, W. E., Bervejillo, J. E., Carriquiry, M., Dobermann, A., et al. (2016). Translating the Sustainable Development Goals into action: A participatory backcasting approach for developing national agricultural transformation pathways. *Global Food Security-Agriculture Policy Economics and Environment, 10*, 71-79, doi:10.1016/j.gfs.2016.08.002.

Kennedy, C. M., Hawthorne, P. L., Miteva, D. A., Baumgarten, L., Sochi, K., Matsumoto, M., et al. (2016). Optimizing land use decision-making to sustain Brazilian agricultural profits, biodiversity and ecosystem services. *Biological Conservation, 204*, 221-230, doi:10.1016/j.biocon.2016.10.039.

Kindred, D. R., Milne, A. E., Webster, R., Marchant, B. P., & Sylvester-Bradley, R. (2015). Exploring the spatial variation in the fertilizer-nitrogen requirement of wheat within fields. *The Journal of Agricultural Science, 153*(1), 25-41, doi:10.1017/S0021859613000919.

Lark, R. M., Milne, A. E., Addiscott, T. M., Goulding, K. W. T., Webster, C. P., & O'Flaherty, S. (2004). Scale- and location-dependent correlation of nitrous oxide emissions with soil properties: an analysis using wavelets. *European Journal of Soil Science, 55*(3), 611-627, doi:10.1111/j.1365-2389.2004.00620.x.

Letourneau, D. K., Jedlicka, J. A., Bothwell, S. G., & Moreno, C. R. (2009). Effects of Natural Enemy Biodiversity on the Suppression of Arthropod Herbivores in Terrestrial Ecosystems. *Annual Review of Ecology, Evolution, and Systematics, 40*(1), 573-592, doi:10.1146/annurev.ecolsys.110308.120320.

Metcalfe, H., Milne, A., & Storkey, J. (2019). (paper in preparation). *Ecology Letters*.

Nelson, E., Mendoza, G., Regetz, J., Polasky, S., Tallis, H., Cameron, D. R., et al. (2009). Modeling multiple ecosystem services, biodiversity conservation, commodity production, and tradeoffs at landscape scales. *Frontiers in Ecology and the Environment, 7*(1), 4-11, doi:10.1890/080023.

O'Farrell, P. J., & Anderson, P. M. L. (2010). Sustainable multifunctional landscapes: a review to implementation. *Current Opinion in Environmental Sustainability, 2*(1-2), 59-65, doi:10.1016/j.cosust.2010.02.005.

Parsons, D. J., & Beest, D. T. (2004). Optimising fungicide applications on winter wheat using genetic algorithms. *Biosystems Engineering, 88*(4), 401-410, doi:10.1016/j.biosystemseng.2004.04.012.

Rawls, W. J. (1983). Estimating soil bulk density from particle size analysis and organic matter content. *Soil Science, 135*(2), 123-125.

Robinson, J. (2003). Future subjunctive: backcasting as social learning. *Futures, 35*(8), 839-856, doi:10.1016/s0016-3287(03)00039-9.

Seppelt, R., Beckmann, M., Ceausu, S., Cord, A. F., Gerstner, K., Gurevitch, J., et al. (2016). Harmonizing Biodiversity Conservation and Productivity in the Context of Increasing Demands on Landscapes. *Bioscience, 66*(10), 890-896, doi:10.1093/biosci/biw004.

Sethi, L. N., Panda, S. N., & Nayak, M. K. (2006). Optimal crop planning and water resources allocation in a coastal groundwater basin, Orissa, India. *Agricultural Water Management, 83*(3), 209-220, doi:10.1016/j.agwat.2005.11.009.

Shibu, M. E., Leffelaar, P. A., van Keulen, H., & Aggarwal, P. K. (2010). LINTUL3, a simulation model for nitrogen-limited situations: Application to rice. *European Journal of Agronomy, 32*(4), 255-271, doi:10.1016/j.eja.2010.01.003.

Storn, R., & Price, K. (1997). Differential evolution - A simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization, 11*(4), 341-359, doi:http://dx.doi.org/10.1023/a:1008202821328.

Teillard, F., Doyen, L., Dross, C., Jiguet, F., & Tichit, M. (2017). Optimal allocations of agricultural intensity reveal win-no loss solutions for food production and biodiversity. *Regional Environmental Change, 17*(5), 1397-1408, doi:10.1007/s10113-016-0947-x.

Tilman, D., Cassman, K. G., Matson, P. A., Naylor, R., & Polasky, S. (2002). Agricultural sustainability and intensive production practices. *Nature, 418*(6898), 671-677, doi:10.1038/nature01014.

Todman, L. C., Coleman, K., Milne, A. E., Gil, J. D. B., Reidsma, P., Schwoob, M.-H., et al. (2019). Multi-objective optimization as a tool to identify possibilities for future agricultural landscapes. *Science of the Total Environment, 687*, 535-545, doi:https://doi.org/10.1016/j.scitotenv.2019.06.070.

van Vliet, M., & Kok, K. (2015). Combining backcasting and exploratory scenarios to develop robust water strategies in face of uncertain futures. *Mitigation and Adaptation Strategies for Global Change, 20*(1), 43-74, doi:10.1007/s11027-013-9479-6.

Verhagen, W., van der Zanden, E. H., Strauch, M., van Teeffelen, A. J. A., & Verburg, P. H. (2018). Optimizing the allocation of agri-environment measures to navigate the trade-offs between ecosystem services, biodiversity and agricultural production. *Environmental Science & Policy, 84*, 186-196, doi:10.1016/j.envsci.2018.03.013.

Westmacott, R. N., & Worthington, T. (2006). *Agricultural Landscapes: 33 Years of Change: Report of a Study Undertaken During 2005 on Behalf of the Countryside Agency's Landscape, Access and Recreation Division*: Countryside Agency.

Wolf, J. (2012). User guide for LINTUL4 and LINTUL4V: Simple generic model for simulation of crop growth under potential, water limited and nitrogen limited conditions (pp. 58). Wageningen UR, Wageningen.

Table 1 Soil properties for the topsoil (0-23cm) of the fields 1 and 2. Here sand has a particle size distribution between 2000-60 µm, silt is between 60-2 µm, and clay is <2 µm.

	Soil - type	Texture			Organic C	pH	Bulk density
		Sand	Silt	Clay			
		%	%	%	%		(g cm^{-3})
Field 1	Clay	9.8	14.3	75.8	2.49	7.6	1.231
Field 2	Sandy loam	68.0	17.9	14.2	0.96	6.0	1.337

Table 2 Time taken for the optimisation to converge and the number of iterations before convergence was achieved. Scenario 1 is a 1x2 grid of cells with no spatial interaction, Scenario 2 is a 1x2 grid of cells with lateral flow (i.e. spatial interaction), Scenario 3 is a 10x10 grid with spatial interaction. Strategy 1 is where the optimisation is applied to individual cells and solutions combined post-hoc, Strategy 2 assumes that management (controls) is applied uniformly across all cells, Strategy 3 optimises the whole grid assuming that management may vary from cell to cell and Strategy 4 is the same as Strategy 3 but with the initial conditions of the control variables partially defined by the results from single cell optimisations.

	Number of control variables	**Number of iterations to convergence**	**Time taken to converge**
Single cell field 1	11	48	32 mins, 16 secs
Single cell field 2	11	70	46 mins, 27 secs
Scenario 1 with Strategy 1	22	85	1 hr, 50 mins
Scenario 1 with Strategy 2	11	30	41 mins, 45 secs
Scenario 2 with Strategy 3	22	77	1 hr, 40 mins
Scenario 2 with Strategy 4	22	24	33 mins, 1 sec
Scenario 3 with Strategy 4	4752	1760	64 days, 8 hrs

Figure Captions

Fig. 1 A schematic of the landscape model showing the processes that are simulated and how they interact.

Fig. 2 (a) A 1km x 1km landscape in East Anglia, UK (b) A map of the elevation of that landscape (c) the course representation of the landscape in the model with each cell (100 m x 100m). The grey areas represent non-agricultural areas (buildings or woods), the coloured squares indicate the rotation that cell is run with. Yellow, light green, dark green and light blue cells are in a six-year rotation of wheat–beans–wheat–barley–wheat–oilseed rape. Each colour starts at a different point in the rotation. The dark blue and orange cells, are in a wheat–wheat–oilseed rape rotation.

Fig. 3 Phosphorus fertilizer only (green), mineral fertilizer and no FYM (blue) and FYM only (orange). Note that, as increases in nitrous oxide emissions are a negative environmental impact, the y-axis shows values increasing downwards resulting in a convex frontier.

Fig. 4 Comparing the results from optimising the landscape in one stage (black open discs) with the two-stage optimisation, where the results from optimising Field 1 are combined with the results from optimising Field 2 (the frontier of the closed discs). The green discs result from simulations where fertilizer P is applied to both fields, the grey discs indicate solutions where fertilizer P is applied in one field and fertilizer-N or FYM is applied in the other. The blue discs indicate solutions where fertilizer-N is applied in both fields and the orange where FYM applied in Field 2 and fertilizer-N in Field 1.

Fig. 5 The optimisation results from the 1x2 cell optimisation with spatial interaction (blue solid discs) compared with the results where there is no interaction (black open discs). In the case where there is spatial interaction nutrients and water flow from Field 1 to Field 2 due to an elevation gradient between the two fields.

Fig. 6 The optimisation results from the 1x2 cell optimisation assuming uniform management across the landscape (red solid discs) compared with the results where the control (fertilizer application) can vary between fields (black open discs).

Fig. 7 Comparing the results from optimising the landscape in (a) one stage with the (b) two-stage optimisation, where the results from optimising Field 1 are combined with the results from optimising Field 2. The black solid discs relate to solutions that comply with the constraint, whereas the red solid discs do not and so the N-leaching limit is exceeded.

Fig. 8 Green P-fertilizer applied to wheat and oilseed rape only, lower levels of mineral N- and P-fertilizer on all crops and lower levels of FYM applied to oilseed rape (blue) and larger levels of mineral fertilizer with FYM (orange). Note that, as increases in nitrous oxide emissions are a negative environmental impact, the y-axis shows values increasing downwards resulting in a convex frontier

601 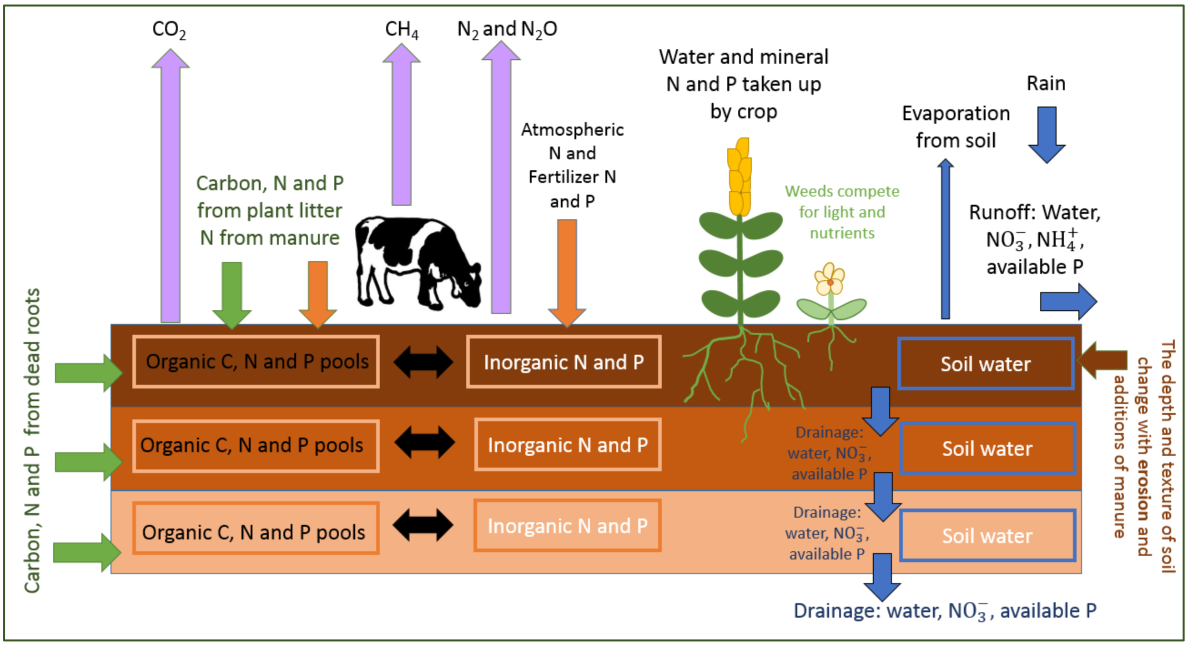

602 Fig. 1 A schematic of the landscape model showing the processes that are simulated and how they

603 interact.

604

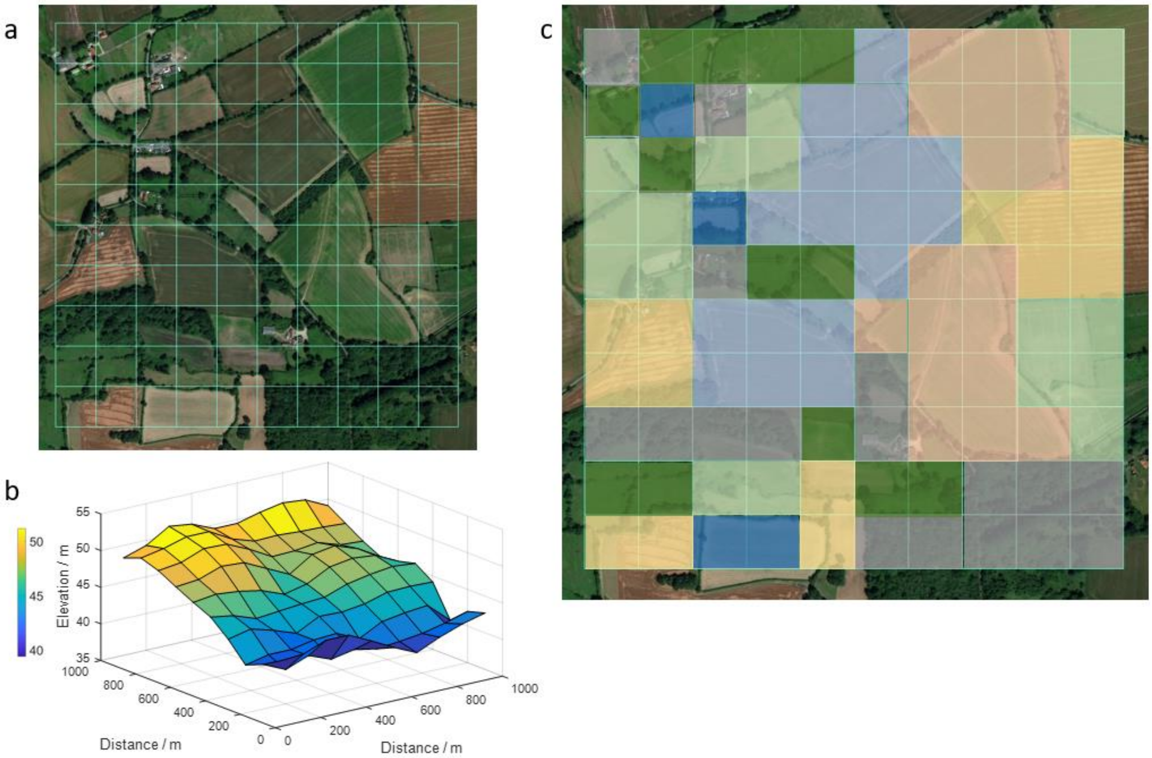

Fig. 2 (a) A 1km x 1km landscape in East Anglia, UK (b) A map of the elevation of that landscape (c) the course representation of the landscape in the model with each cell (100 m x 100m). The grey areas represent non-agricultural areas (buildings or woods), the coloured squares indicate the rotation that cell is run with. Yellow, light green, dark green and light blue cells are in a six-year rotation of wheat–beans–wheat–barley–wheat–oilseed rape. Each colour starts at a different point in the rotation. The dark blue and orange cells, are in a wheat–wheat–oilseed rape rotation.

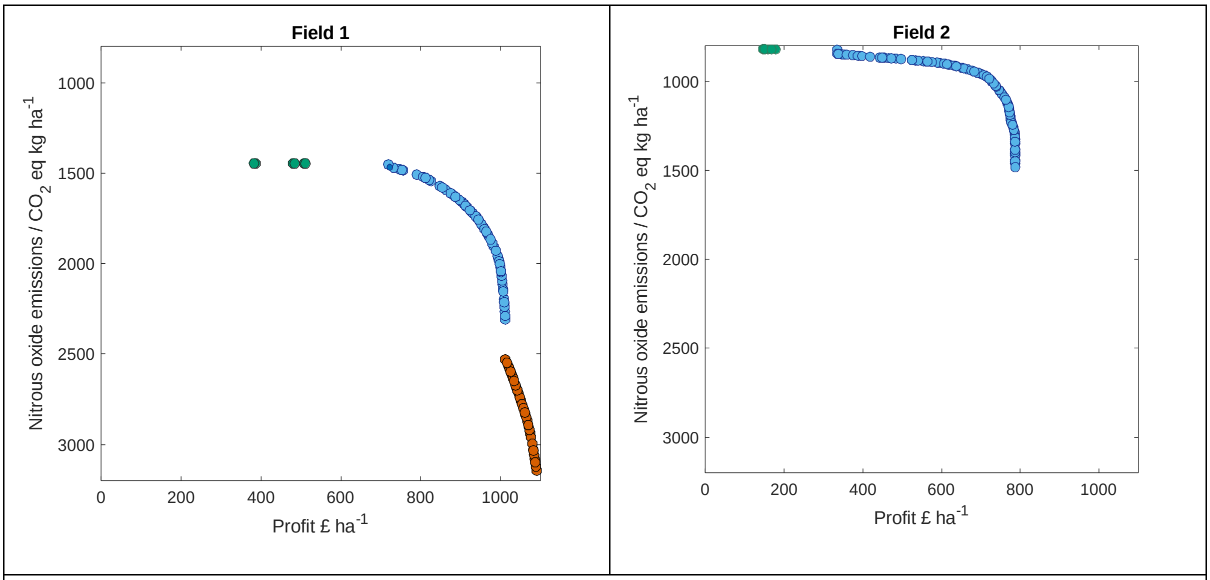

Fig. 3 Green P fertilizer only, mineral fertilizer FYM (blue) and FYM only (orange). Note that, as increases in nitrous oxide emissions are a negative environmental impact, the y-axis shows values increasing downwards resulting in a convex frontier.

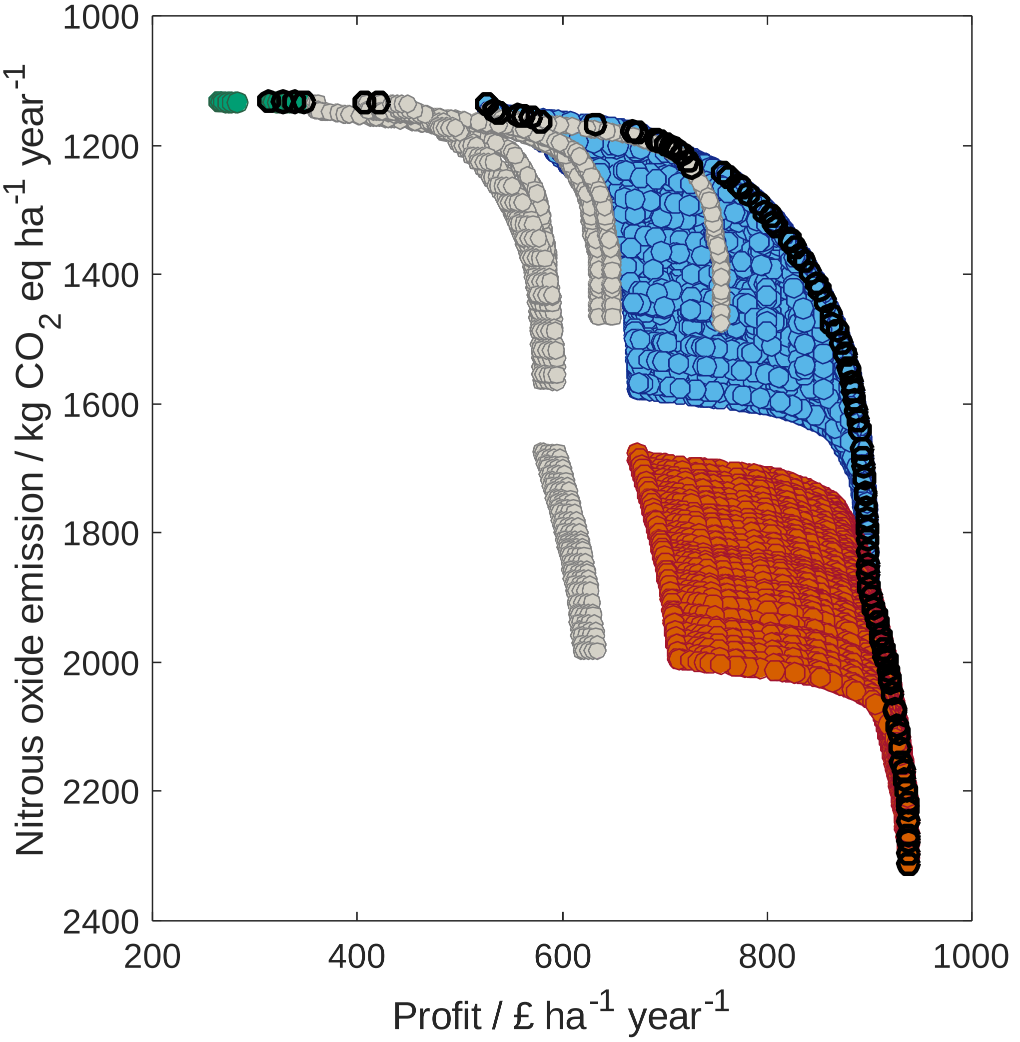

Fig. 4 Comparing the results from optimising the landscape in one stage (black open discs) with the two-stage optimisation, where the results from optimising Field 1 are combined with the results from optimising Field 2 (the frontier of the closed discs). The green discs result from simulations where fertilizer P is applied to both fields, the grey discs indicate solutions where fertilizer P is applied in one field and fertilizer-N or FYM is applied in the other. The blue discs indicate solutions where fertilizer-N is applied in both fields and the orange where FYM applied in Field 2 and fertilizer-N in Field 1.

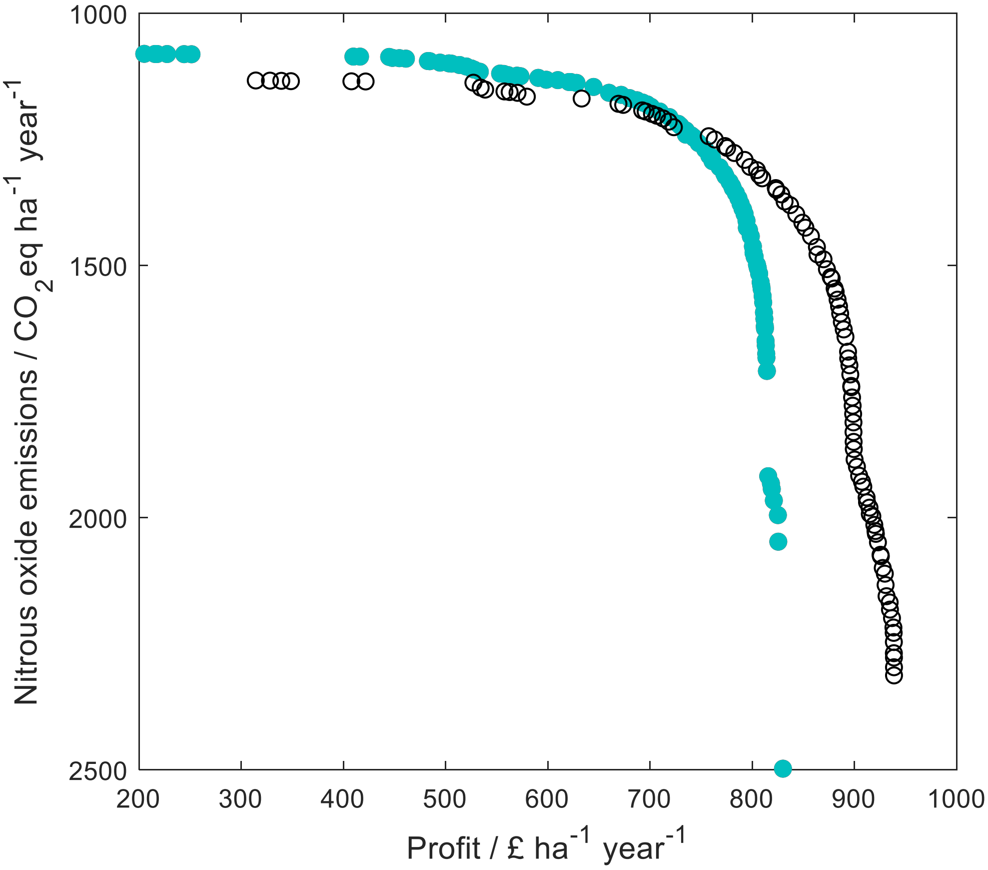

Fig. 5 The optimisation results from the 1x2 cell optimisation with spatial interaction (blue solid discs) compared with the results where there is no interaction (black open discs). In the case where there is spatial interaction nutrients and water flow from Field 1 to Field 2 due to an elevation gradient between the two fields.

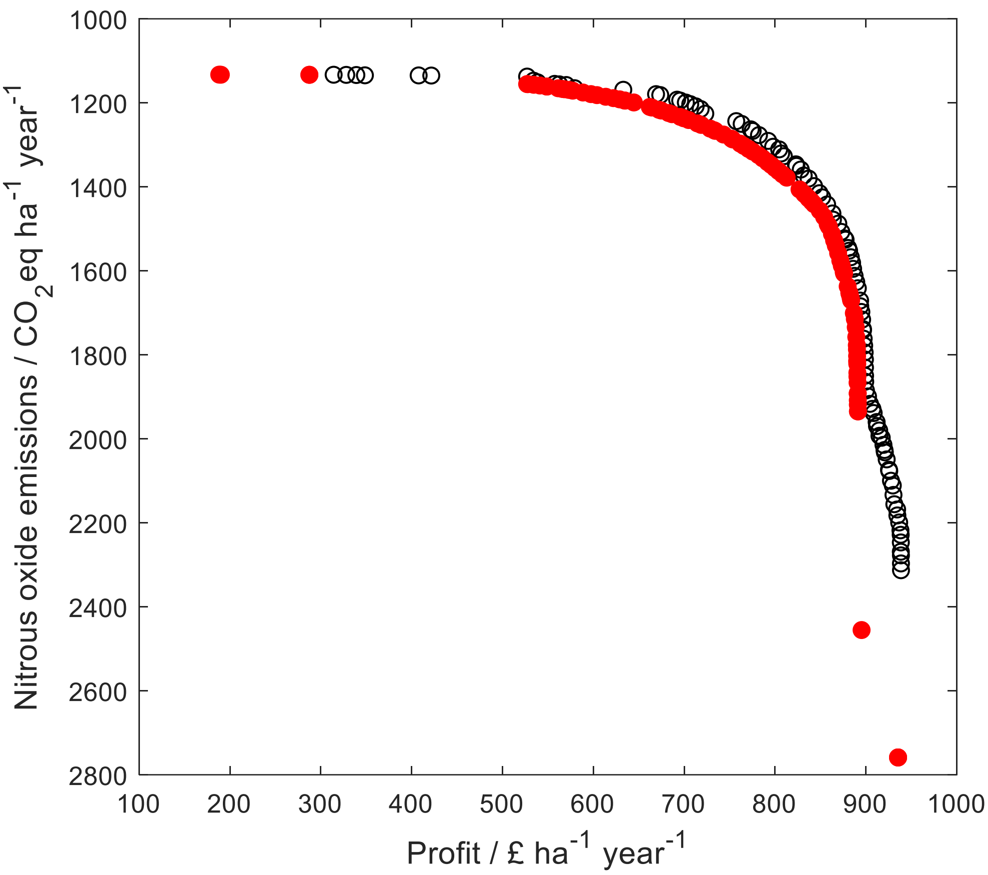

Fig. 6 The optimisation results from the 1x2 cell optimisation assuming uniform management across the landscape (red solid discs) compared with the results where the control (fertilizer application) can vary between fields (black open discs).

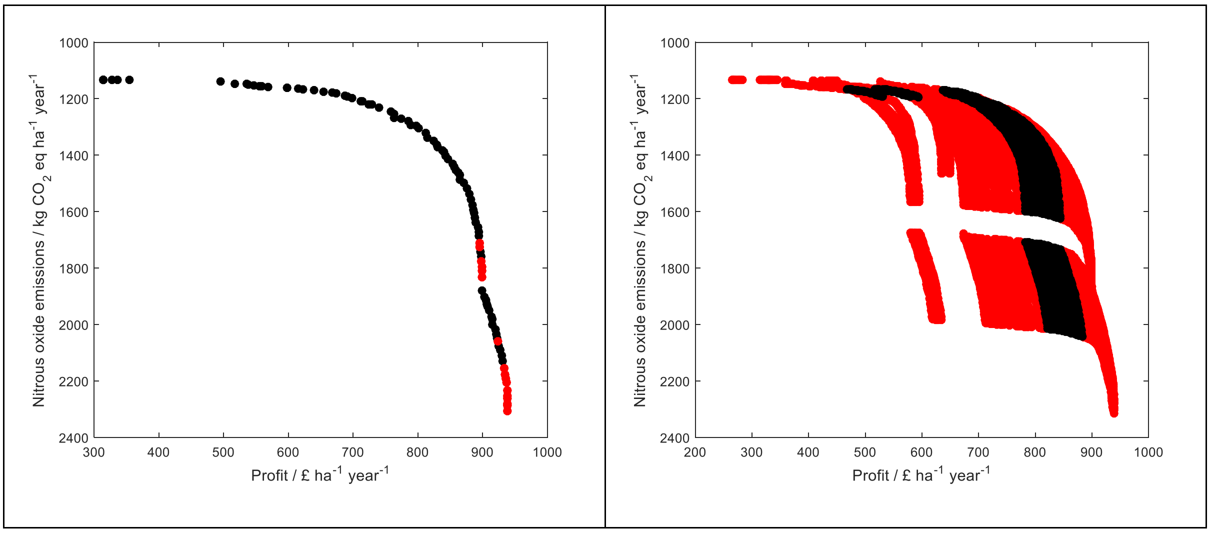

Fig. 7 Comparing the results from optimising the landscape in (a) one stage with the (b) two-stage optimisation, where the results from optimising Field 1 are combined with the results from optimising Field 2. The black solid discs relate to solutions that comply with the constraint, whereas the red solid discs do not and so the N-leaching limit is exceeded.

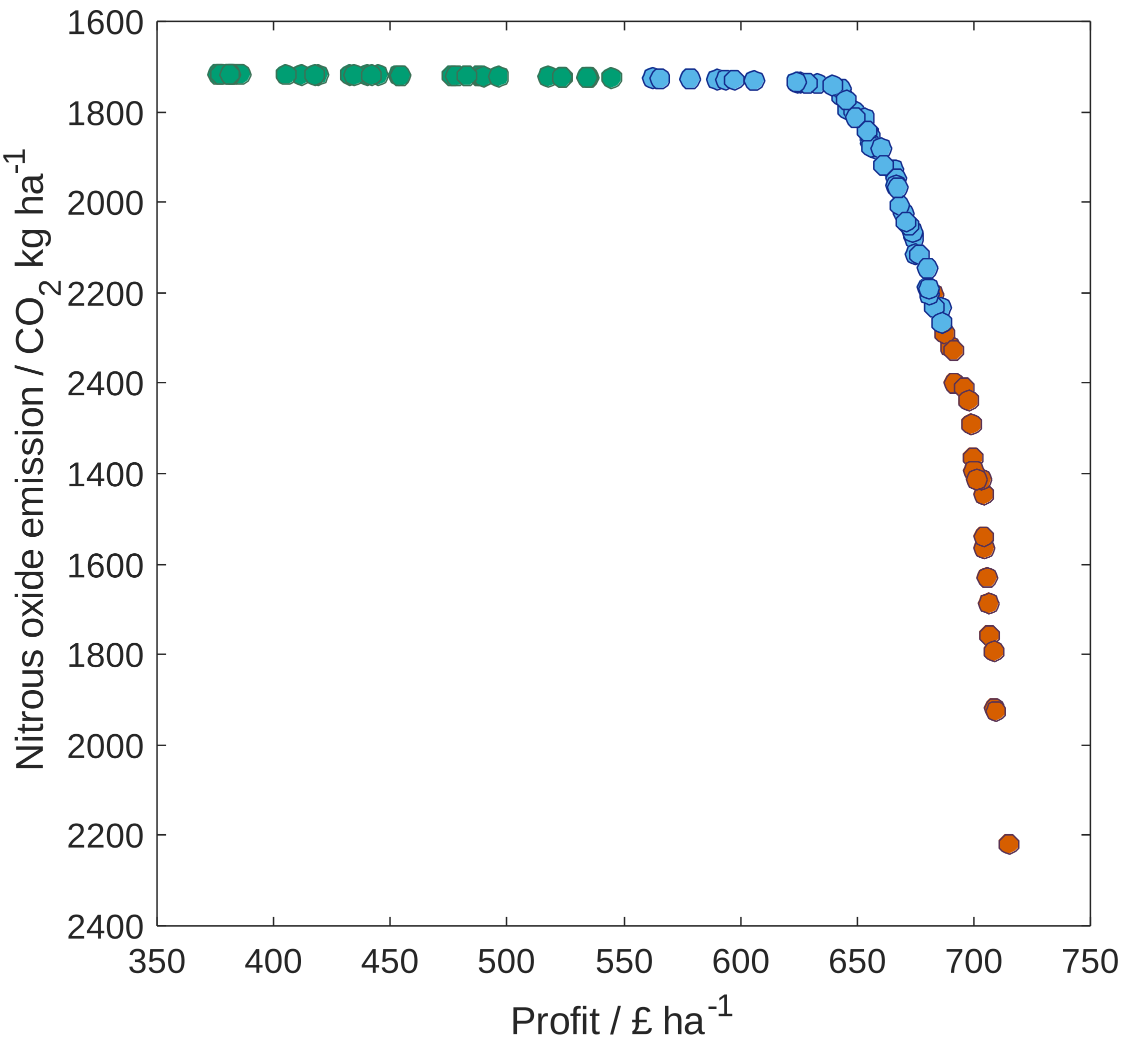

Fig. 8 Green P-fertilizer applied to wheat and oilseed rape only, lower levels of mineral N- and P-fertilizer on all crops and lower levels of FYM applied to oilseed rape (blue) and larger levels of mineral fertilizer with FYM (orange). Note that, as increases in nitrous oxide emissions are a negative environmental impact, the y-axis shows values increasing downwards resulting in a convex frontier.