Received: 10 July 2023 | Accepted: 22 March 2024

DOI: 10.1002/aepp.13448

SUBMITTED ARTICLE

Agricultural land use modeling and climate change adaptation: A reinforcement learning approach

Christian Stetter | **Robert Huber** | **Robert Finger**

Agricultural Economics and Policy Group, ETH Zurich, Zürich, Switzerland

Correspondence
Christian Stetter, Sonneggstrasse 33, 8092 Zürich, Switzerland.
Email: cstetter@ethz.ch

Editor in charge: Ashok Mishra

Abstract

This paper provides a novel approach to integrate farmers' behavior in spatially explicit agricultural land use modeling to investigate climate change adaptation strategies. More specifically, we develop and apply a computationally efficient machine learning approach based on reinforcement learning to simulate the adoption of agroforestry practices. Using data from an economic experiment with crop farmers in Southeast Germany, our results show that a change in climate, market, and policy conditions shifts the spatial distribution of the uptake of agroforestry systems. Our modeling approach can be used to advance currently used models for ex ante policy analysis by upscaling existing knowledge about farmers behavioral characteristics and combine it with spatially explicit environmental and farm structural data. The approach presents a potential solution for researchers who aim to upscale information, potentially enriching and complementing existing land use modeling approaches.

KEYWORDS

agroforestry, climate change adaptation, land use modeling, machine learning, reinforcement learning

JEL CLASSIFICATION
Q15, Q18, C31, D91

Understanding the impacts of climate change on land use and economic outcomes requires knowledge of farmers' adaptation responses to changing climatic conditions (Burke & Emerick, 2015). Agricultural land use models are a key instrument for investigating farmers' land use decisions under climate change and to assess the potential of climate policies (Fei & McCarl, 2023; Troost & Berger, 2014). They are of specific relevance to model climate change adaptation measures and increasing incidences of extreme weather events (e.g., Mendelsohn et al., 1994; Nelson et al., 2014; Ortiz-Bobea et al., 2021). These adaptation measures are highly relevant to food security, economic viability, incomes, and the environmental impacts of the farming sector (e.g., Carter et al., 2018; Fisher et al., 2012; Wheeler et al., 2013). Thus, ex ante modeling assessments allow us to investigate land use decisions under changing climatic, market, and policy conditions are of high policy relevance (Fei & McCarl, 2023). Existing modeling approaches, however, deal with key trade-offs, especially between the consideration of behavioral factors, spatial resolution, as well as the data demand and requirement of computational resources.

In this study, we address the question of how information about farmers' behavioral characteristics can be linked to spatially explicit environmental data for ex ante assessments of climate and policy impacts on agricultural land use. More specifically, the aim of this study is to present a computationally efficient machine learning approach based on reinforcement learning to simulate—in a spatially explicit manner—the potential uptake of agroforestry practices as a strategy to adapt to climate change.

Various studies provide ex post evaluations of how weather and climate change have influenced agriculture, alongside examining the extent to which farms have implemented adaptation measures (e.g., Auffhammer, 2018; Lobell, 2014; Mérel & Gammans, 2021). However, climate-induced land use adjustments will involve structural shifts, such as the switch to novel land use types, practices, and technologies that are currently not in use. Thus, projections of past land use change based on observations of previous decisions alone have limited value for assessing farmer behavior or introducing new climate policies (Antle & Capalbo, 2001; Fei & McCarl, 2023; Troost & Berger, 2014). In this context, ex ante assessments are an important complementary approach to ex post evaluations that inform policies for climate change induced land use adaptation in agriculture and for assessing the effects of implemented policies.

A wide range of modeling approaches are available, and they range from the field, farm, and regional to the global level (see e.g., Arneth et al., 2014; El Benni et al., 2023; Lehmann et al., 2013; Nelson et al., 2014; Troost & Berger, 2014). However, these modeling approaches are fraught with important trade-offs (Rindfuss et al., 2004), especially between the in-depth consideration of behavioral factors, for example, from experimental economic approaches (Huber et al., 2018; Reidsma et al., 2018), spatial resolution of models (Fezzi & Bateman, 2011; Nolte, 2020), and data demand and computational resources (Shang et al., 2023; Troost et al., 2022). For example, improving models by adding spatially explicit behavioral information has high policy value (Verburg et al., 2019). However, spatially explicit behavioral models may reinforce the need for data and computational resources (Brown et al., 2017). A promising new pathway is the use of modern machine learning algorithms (Wang et al., 2022). Such algorithms have successfully been utilized in the field of agricultural economics, for example, to study the impacts of extreme weather events (e.g., Webber et al., 2020), yield predictions (e.g., Schmidt et al., 2022), agricultural policy evaluation (e.g., Stetter et al., 2022), or food security prediction (McBride et al., 2022). Although applied machine learning rapidly evolving, the deep integration of big data and machine learning with economic research is still limited (e.g., Athey, 2017; Wang et al., 2022). In particular, spatially explicit models of agricultural land use change under climate change lack computationally efficient algorithms to account for behavioral factors such as farmers' preferences.

To address this gap, we provide a novel land use modeling approach based on reinforcement learning using contextual multi-armed bandits. We integrate farmers' stated preferences for agroforestry practices, revealed through a discrete choice experiment on crop farms located in Southeast Germany with farm-specific data on their short- and long-term weather history. Based on these data, we provide a spatially explicit ex ante assessment of the adoption of agroforestry practices under climate, market, and policy scenarios. The modeling approach is inspired by contextual multi-armed bandits applied in recommender systems,[1] the key principles of which we incorporate into our land use modeling approach. Contextual multi-armed bandits are particularly beneficial in our context because of their ability to learn rules, that map from observable characteristics of an individual or location (e.g., the farm) to an action (e.g., the adoption of agroforestry).

This study contributes to the literature in two ways. First, we introduce a novel modeling approach that combines empirical behavioral economic research approaches, including the use of economic experiments, with reinforcement learning. Our approach presents a potential solution for researchers who aim to upscale information, potentially enriching and complementing existing modeling approaches. Second, we provide site-specific recommendations for the adoption of agroforestry practices, with a prescriptive character that can inform farmers' decision-making process in the context of novel land use types that are more resilient to climate change and currently not commonly used by farmers. This approach is relevant for identifying and implementing sustainable land use practices that can alleviate the adverse effects of climate change on agriculture and food production.

Our results exemplify a successful integration of behavioral and environmental data in a spatially explicit manner. Thus, our approach extends the land-use modeling toolbox. The simulations show that a change in climate, market, and policy conditions leads to shifts in the spatial distribution of the uptake of agroforestry systems. This might provide relevant information for farmers, farm advisors, or policymakers seeking to foster the adaptation to climate change.

The remainder of the article is structured as follows. Background Section 1 provides contextual information regarding agroforestry as a climate change adaptation strategy. Section 2 presents the conceptual framework of this study. Section 3 describes the data used, and Section 4 explains the modeling implementation. Section 5 presents the results of our analysis, Section 6 discusses the findings, and Section 7 concludes.

BACKGROUND: AGROFORESTRY AS CLIMATE CHANGE ADAPTATION STRATEGY

Land use and related adjustments play a critical role in adapting to and alleviating the adverse effects of climate change (Pielke, 2005). These adjustments involve the cultivation of alternative crops and varieties along with adjustments at the intensive margin including the use of water, fertilizers or pesticides (e.g., Cabral et al., 2023; Graveline & Merel, 2014; Wimmer et al., 2023). The Intergovernmental Panel on Climate Change recommends the cultivation of agroforestry systems as an effective climate change adaptation measure (IPCC, 2022).

Agroforestry offers a solution for managing agricultural risks associated with climate variability and extreme events. Incorporating trees into agroforestry systems alleviates climate fluctuations, improves microclimates, alters wind patterns and temperatures, and regulates water balance. Furthermore, agroforestry enhances soil health, fertility, and biodiversity, thereby bolstering agroecosystem resilience to climate change (van Noordwijk et al., 2021). It also generates multiple additional income streams and helps alleviate the adverse effects of extreme weather events on crop productivity and profitability. Overall, incorporating woody crops into agricultural landscapes can help build resilience and adaptation to the impacts of climate change while also providing multiple benefits to the environment and society (van Noordwijk et al., 2021; Wolz et al., 2018).

Here, we focus on the integration of fast-growing trees such as poplar (*Populus* spp.), willow (*Salix* spp.), and black locusts (*Robinia* spp.) in arable lands. These trees are usually harvested (i.e., cut) every 2–5 years for bioenergy production. They are characterized by rapid re-growth of shoots from stumps. The total lifetime of such trees is about 20 years, after which they are usually converted back into arable cropping systems (Rödl, 2017). We specifically examined two types of agroforestry systems: short-rotation coppice and alley-cropping. Fast-growing trees are cultivated *instead* of traditional arable crops in short-rotation coppice systems (Rödl, 2017). By contrast, in alley-cropping systems rows of trees are planted *alongside* agricultural crops in strips or alleys. Alley spaces between tree rows are used for cultivating crops. The trees are typically planted in such a way that competition with other crops is avoided (Wolz et al., 2018).

CONCEPTUAL MODELING FRAMEWORK

Land use is a consequence of the interplay between many drivers and processes including natural, institutional, economic, and behavioral factors (Meyfroidt et al., 2018; Wang et al., 2022). Thus, simulating land use changes under climate change depends on the integration of these factors. This is of specific importance in the context of agricultural land use change which depends on spatially diverse farming conditions, diversified farming systems, and the decision-making of heterogenous farmers (Huber et al., 2018).

The proposed machine learning approach integrates spatially explicit information on farmers' land use preferences for agroforestry practices with economic and environmental information for upscaling land use change in a computationally efficient manner (Figure 1). This model helps provide site-specific recommendations.

In the following sections, we describe the modeling approach in more detail. We identify farmers' preferences and integrate them with environmental data in order to predict agricultural land-use changes in space, that is, across heterogeneous farms and farming conditions under climate and policy changes.

Random utility theory and farmers' preferences

The behavioral model that guides individual farmers' decision-making processes is assumed to follow random utility maximization (Lancaster, 1966; McFadden, 1973). This flexible behavioral framework explains how individuals choose among alternatives in the presence of uncertainty. Regarding land use in the context of climate change, farmers must choose among a set of alternative land uses. Each land user obtains a certain level of indirect utility from each land-use alternative. In decision situation t, they select alternative i if and only if:

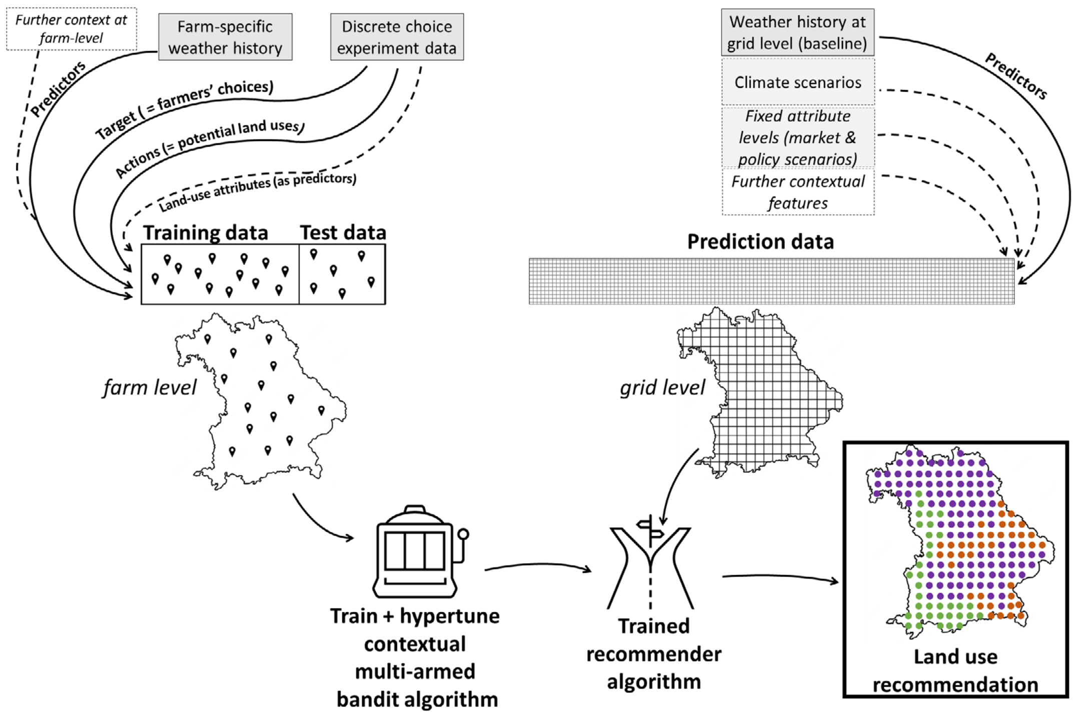

FIGURE 1 Conceptual modeling framework. To optimize land use recommendations for farmers, data from a discrete choice experiment deliver information on farmers' choices (stated preferences), potential land uses (i.e., their action space), and information on land use attributes. Local weather history is used to describe farmers' context. However, the approach also allows to include further contextual information. The farm-level data are then used to train and hypertune a contextual multi-armed bandit algorithm, which yields a trained land-use recommender algorithm that maps context to land-use decisions. The trained algorithm is subsequently supplied with raster-level contextual data, enabling it to generate land-use recommendations for each raster unit. The contextual factors influencing land-use decisions (weather history and land use attributes) can vary depending on the specific scenario.

$$U_{it} > U_{jt},\ j \neq i. \tag{1}$$

As a standard, it is assumed that farmers' utility from a land-use alternative varies with a set of decision-relevant characteristics. In the case of the here used experiment presented by Stetter and Sauer (2022), these are contribution margin[2] contribution margin variability, minimum useful lifetime, agri-environmental payments, and eligibility as an ecological priority area. Following Nerlove (1958) and Ramsey et al. (2021), we assume that farmers' subjective utility also depends on their experiences with short- and long-term weather trends, that is, the local weather history at the time of the planting decision. Short- and long-term weather trends might affect land use choices. The direct link between local weather history and individual decisions comes from the fact that past weather causes various types of subjectively experienced uncertainties (compare Stetter & Sauer, 2022). This theoretical framework builds upon a growing body of research that highlights the significant link between historical weather events and economic decision-making. For example, Ramsey et al. (2021) and Wimmer et al. (2023) demonstrate that land-use decisions are shaped by a combination of local short- and long-term weather histories. Ji and Cobourn (2021) further substantiate this finding, indicating that economic decision-making is driven by both short-term and long-term experienced weather from the past, influencing agents' subjective expectations regarding future climate conditions. Petersen-Rockney (2022) finds that experienced weather holds greater relevance for climate change adaptation compared to farmers' perceptions thereof.

Discrete choice experiments are a popular method that can be used to empirically elicit individuals' preferences based on random utility theory (Louviere et al., 2000). Using the results of such an experiment in a land use recommendation model allows for the explicit inclusion of real-world, empirically measured farmer behavior without relying on oversimplifying assumptions.

Data-driven recommender system for optimal land use under climate change

Recommender systems generate recommendations for users based on their interests and needs and play a crucial role in digital companies, leveraging big data to improve performance (Buhmann et al., 2011). In recent years, the multi-armed bandit framework has gained significant attention in recommender systems (Bouneffouf & Rish, 2019). Multi-armed bandits fall under the broader category of reinforcement learning, which is a branch of machine learning dedicated to understanding how intelligent agents make optimal decisions in dynamic environments (Pathak et al., 2021).[3]

Multi-armed bandits are a class of problems for adaptive decision-making under uncertainty, where an agent must choose between multiple options, each with an unknown reward distribution. The agent's goal is to maximize the total reward obtained over a sequence of choices, while balancing the trade-off between exploration (trying out different options to learn about their rewards) and exploitation (choosing the action that appears to be the best based on current knowledge) (Pathak et al., 2021). We frame the agent's land use allocation problem as a contextual multi-armed bandit problem, which is a variant of the classical multi-armed bandits problem. In the contextual multi-armed bandit problem, the reward distribution depends not only on land users' decisions but also on the context associated with it (Speekenbrink, 2022).

At the fundamental level, within multi-armed bandit problems, an agent faces an assignment problem with limited information (i.e., uncertainty). Given the novel nature of climate-resistant land uses (i.e., agroforestry), there is a lot of uncertainty surrounding the attributes of these novel land use options, which are usually unknown or unobserved in the real-world (compare e.g., Lu et al., 2010). However, the local weather history can be measured. Hence, in our model, the task is to map weather history (i.e., context) to an action (i.e., the decision to choose either alley-cropping, short-rotation coppice, or arable crop farming), such that the expected cumulative expected reward (i.e., utility) across the entire sample (and thus all land users) is maximized.

Mathematically, let a_t be the chosen action (i.e., choosing a land use type) in decision situation t and c_t be the context or features associated with each action (i.e., weather history) in t. Let $r(a,c)$ denote the distribution of rewards obtained by performing action a in context c. The expected cumulative reward function over contexts for action a is defined as (compare Silva et al., 2022; Vermorel & Mohri, 2005):

$$Q(a,c) = E\left[\sum_{t=1}^{N} r(a_t, c_t)\right], \tag{2}$$

where N is the number of observed decision situations (e.g., from a discrete choice experiment). In other words, the contextual multi-armed bandit algorithm aims to learn a mapping from context[4] c_t to action a_t, such that the expected cumulative reward is maximized across all observations. The choice of action in each decision situation depends on the context observed at that time, and the algorithm updates its estimates of the expected rewards of each action given the observed contexts and rewards over time (i.e., learns from each decision situation, see also Figure S1).[5] By adjusting the contextual information provided to the agent, various policy and climate-change scenarios can be simulated.

Finally, the agent's goal is to determine the action a^* that maximizes the expected cumulative reward for each context c (Vermorel & Mohri, 2005):

$$a^*(c) = \arg\max_a Q(a,c). \tag{3}$$

The mapping function $a^*(c)$ is learned from observations by repeatedly selecting actions, observing the corresponding rewards and contexts, and updating the estimates of $Q(a,c)$.

CASE STUDY AND DATA

Our case study region Bavaria is a German federal state located in the southeast of the country (Figure 2a). It belongs to the core regions of agricultural production in Germany and the European Union (Stetter et al., 2023). In 2022, with 78,000 farms managing 3.1 million ha, crop farming occupied 65% of the agricultural land (Bavarian State Office for Statistics, 2023). The natural conditions in Bavaria vary, with farming occurring across a gradient of elevation from 100 to 1600 m and a range of annual average temperatures between 3 and 10°C. Annual precipitation levels also vary from 470 to 1592 mm (LfU, 2019). Climate change implies more extreme weather events that will threaten the viability of Bavarian farms (Bavarian State Ministry for the Environment and Consumer Protection, 2021). Notably, the investigated agroforestry systems, namely, alley-cropping and short-rotation coppice are not (yet) widely cultivated in the case study region.

Choice experiment and historical weather data

The data used to train and test our model were obtained from an online discrete choice experiment conducted with Bavarian crop farmers (Stetter & Sauer, 2022). Each participant was confronted with a total of 12 hypothetical choice situations. The decision situations were defined such that the participants could choose among three land use options: alley-cropping, short-rotation coppice, or business-as-usual arable crop farming.[6] The attributes of the agroforestry options varied across choice sets in five dimensions: contribution margin, contribution margin variability, minimum useful lifetime, agri-environmental payments, and eligibility as ecological priority area. The ranges of the attribute levels were designed to reflect realistic choice situations that farmers could relate to. A discussion on this can be found in the supplementary material (section S.2) of the paper. All land use options and their characteristics were described to the respondents before starting the experiment.

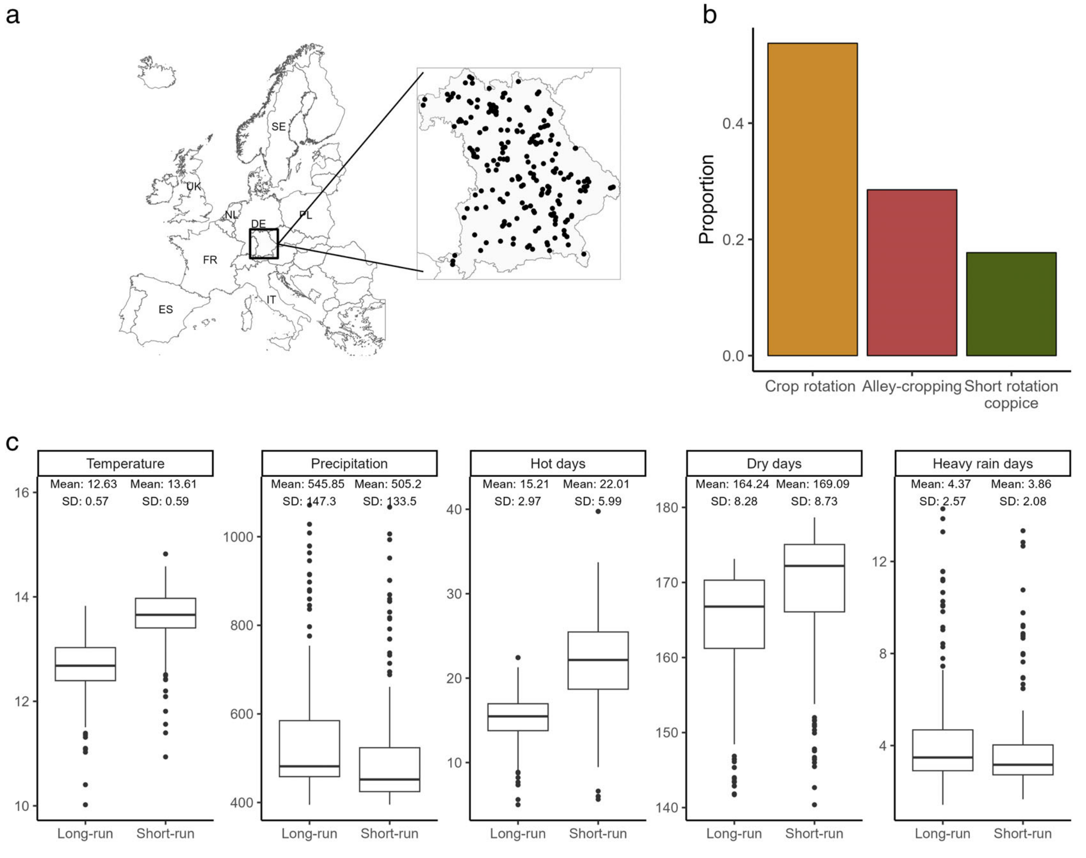

FIGURE 2 Case study and modeling data description. (a) The case study region Bavaria lies in the center of Europe. In total, the preferences of 198 farmers spread across Bavaria were used for the analysis. (b) Shares of land use choices made by the respondents in the discrete choice experiment. (c) Description of weather history data, long-run: years $t-4$ to $t-10$, short-run: years $t-1$ to $t-3$.

The experiment was conducted online in October 2020 and targeted respondents from a diverse pool of farmers sourced through agriEXPERTS, an agricultural market research platform associated with Deutscher Landwirtschaftsverlag (dlv), a specialist publishing house for agriculture. Further information on the discrete choice experiment can be found in Stetter and Sauer (2022).[7]

In total, 198 farmers scattered across Bavaria provided complete responses (see Figure 2a). The participating farmers chose the business-as-usual arable crop farming alternative most often among the three options (1277 times), followed by alley-cropping (678 times) and short-rotation coppice (421 times) (Figure 2b). The data provided in total 7128 farmer-action observations. The target variable takes a value of 1 if a land-use alternative is chosen, and 0 otherwise. Farmers in the sample reported experiencing negative consequences of climate change-related extreme weather events, particularly yield and quality losses (compare Stetter & Sauer, 2022).

Information on the contribution margin (which lies in the range €400–€800 per ha of arable land) and agri-environmental payments for agroforestry (€0–€200 per ha) from the discrete choice are considered in the contextual information to generate market and policy scenarios, respectively (see Section 4.3).

Farmers' stated land use choices were spatially matched with gridded daily weather data at the location of their farm from the European Climate Assessment & Dataset project (Cornes et al., 2018) via zip codes. To accurately depict the local weather history specific to the farms, five commonly used weather indicators were chosen. These indicators encompass the average temperature, total precipitation, number of dry days, number of hot days, and number of heavy rain days observed throughout the local growing season, which spans from March to October.[8] Drawing upon the definitions by ETCCDI (2018) and DWD (2022), the term "dry days" refers to days with precipitation levels below 1 mm, while "hot days" are characterized by maximum temperatures surpassing 30°C. Heavy rain days were classified as those on which precipitation exceeded 20 mm.

In accordance with the theoretical framework discussed in Section 2, farmers rely on historical weather patterns when selecting land-use types, which can be categorized into short- and longer-term patterns. To capture this difference, short- and long-term weather variables with distinct lag structures were defined. Following Ramsey et al. (2021) and Wimmer et al. (2023), the short-term weather patterns for the five indicators were established based on the average values from years $t-1$ to $t-3$ (i.e., 2017–2019), representing the more recent past. Conversely, longer-term weather patterns were derived from the average values of the years $t-4$ to $t-10$ (i.e., 2010–2016), representing a more distant historical period. The weather histories of the farms are summarized in Figure 2c.

Prediction data

Gridded weather history

To build our prediction dataset to project local farmers' decisions in space, we divided the case study region into 1 km^2 regular grid cells ($N = 61{,}205$) and described them based on their individual weather history. Each grid cell was regarded as an individual decision-making unit. Raster cells without cropland are excluded from the analysis. The meteorological data come from the German Meteorological Service (Deutscher Wetterdienst, DWD). Weather history indicators were based on open access daily gridded weather data (DWD Climate Data Center [CDC], 2023). These data are created by interpolating weather measurements from more than 1300 stations (Razafimaharo et al., 2020). Further information on this data set is available in Rauthe et al. (2013) and Razafimaharo et al. (2020). Figure S3 provides an overview of the indicators used for the corresponding land-use predictions.

Climate projections

The data used to assess the impact of climate change on the cultivation of agroforestry in our case study region stems from bias-adjusted climate change scenarios from the Bavarian Climate Projection Ensemble provided by the Bavarian State Office for the Environment (LfU, 2020; Zier et al., 2020, for details). The data were provided on a 5 km × 5 km grid, which yields 5183 locations. It represents daily data aggregated for the growing season as previously performed.

To showcase the usefulness and flexibility of our suggested modeling approach against the background of climate change, we focused on two different greenhouse gas Representative

Concentration Pathways (RCPs),[9] RCP 8.5, and RCP 4.5, based on the EC-Earth global climate model system implemented by the climate-limited area modeling community (Will & Hense, 2008). We assessed land use allocation for 2020, 2030 and 2050. Figure S4 provides an overview of the input data used for the climate change scenarios.

To assess the sensitivity of the obtained results, additional regional climate projections were tested for RCP 2.6, RCP 4.5, and RCP 8.5 in addition to the selected model. Further information is provided in the Table S1.

MODEL IMPLEMENTATION

We employed three widely adopted and well-established contextual multi-armed bandit algorithms (Silva et al., 2022), that is, LinUCB, decision tree UCB (upper confidence bound) and k-nearest neighbor UCB, to develop our land use allocation model. To identify the best-performing recommendation model, each of these algorithms underwent rigorous hyper-parameter tuning. Once the best algorithm was chosen based on predefined criteria, it was employed to investigate the land use implications of different climate, market, and policy scenarios.

Three contextual multi-armed bandit algorithms

All algorithms presented herein consist of three components. First, parametric or non-parametric models predict the expected rewards from each land-use type for any observation. Second, the exploration factor determines the mechanism by which the algorithm explores different actions to gather information regarding potential rewards. The uncertainty component refers to a lack of knowledge or uncertainty regarding the true rewards associated with each action.

LinUCB algorithm

The linUCB algorithm (Li et al., 2010) is a parametric regression-based technique. First, following Strong et al. (2021), the algorithm trains a ridge regression (Hoerl & Kennard, 1970) for each land use type that models the expected reward using a linear relationship between contextual features (X_a)[10] and observed rewards (r) for each land use type a:

$$\hat{r}_a = \boldsymbol{\theta}^T X_a, \tag{4}$$

where $\boldsymbol{\theta}^T$ denotes the parameter vector. Based on this, the algorithm calculates the UCB for each land-use choice for each new observation as follows:

$$\mathrm{UCB}_a = \hat{r}_a + \alpha\sqrt{\left(X_a^T X_a + \lambda \cdot I_d\right)^{-1} x_c}, \tag{5}$$

where α is the exploration factor, λ is the ridge regression strength, and X_a is the design matrix for land use a, which contains the features of actions that have already been explored. x_c describes the feature vector associated with context c. The exploration factor controls the balance between exploration and exploitation. A higher exploration factor encourages more exploration of different land uses, whereas a lower value favors exploitation of the arms that have shown higher estimated rewards. The regularization parameter λ determines the amount of shrinkage applied, with larger values leading to greater shrinkage (Hoerl & Kennard, 1970). $\sqrt{\left(X_a^T X_a + \lambda \cdot I_d\right)^{-1} x_c}$ denotes the predictive precision of the expected reward. The LinUCB algorithm selects the action (i.e., land use) with the highest upper confidence bound value (Li et al., 2010). This approach allows the algorithm to explore actions with high potential rewards and exploit actions that have performed well in the past. At each new observation during training, the coefficients of Equation (4) are updated, restarting the process described above.

Decision tree UCB algorithm

In this approach, a decision tree algorithm is combined with the UCB1 algorithm (Strong et al., 2021). As before, the UCB consists of three parts: the predicted reward for an observation, the exploration factor, and uncertainty. This can be expressed as follows:

$$\mathrm{UCB}_a = \widehat{r}_a + \alpha \times \sqrt{\frac{2 \times \log(N)}{n_a}}, \tag{6}$$

where N is the total number of sample rounds, and n_a is the number of times a land use alternative is selected up to the current time step. In this approach, the Chernoff–Hoeffding bound is used to determine the prediction intervals for estimating the true reward. This decreases as the number of samples from an action (n_a) increases (Auer et al., 2002). By choosing the action with the highest UCB for each training observation, the algorithm effectively balances exploration and exploitation. This allows for efficient learning and decision-making when the goal is to maximize the cumulative reward over samples (McNellis et al., 2017).

As for the prediction of the expected reward $\widehat{r}_a$, a decision tree for each land-use type is trained based on the contextual history X. The decision tree partitions the context space into neighborhoods, and for each leaf node, a list of rewards is maintained. The tree-building process involves recursively splitting the data based on contextual variables, creating a tree structure in which each leaf node (i.e., the terminal node) corresponds to a specific subset of the feature space and provides a prediction of the target variable associated with that region. The predictions at the leaf nodes are calculated as the mean of the observed rewards associated with that leaf (see e.g., Hastie et al., 2009).

As before, for each new observation during training, the algorithm re-starts the process described above. Other than before, decision trees are non-parametric. However, there are many hyperparameters that can be adjusted in order to improve the prediction (compare Section 4.2).

K-nearest neighbors UCB algorithm

This method is similar to the previous method. However, it differs in that it uses k-nearest neighbors regression instead of a decision tree to predict the expected rewards $\widehat{r}_a$ (Strong et al., 2021). The idea behind this is to estimate the reward of an action, examine the contextual variables and to determine the k nearest (most similar) observations based on these variables. The value of k is a predetermined number to be chosen. Subsequently, the rewards of those k nearest observations are averaged. The average of each k-nearest neighbor partition becomes the estimate of the reward for land use i (see e.g., Hastie et al., 2009). The nearest neighbors are determined based on the Euclidean distance. As before, each observation updates the algorithm and re-starts the process.

Hyperparameter tuning and model selection

To select a recommender model, we evaluated different versions of the proposed algorithms. To do so, we split our modeling data into training and test data. Following Dobbin and Simon (2011), we allocated two-thirds of the cases to training. Using the training data, we conducted a hyperparameter grid search across multiple hyperparameter specifications for each algorithm, yielding a comparison of 136 candidate models (Table S2). Using the test data, we evaluated the performance of each model specification based on five criteria. First, we evaluated cumulative expected rewards as our primary target measure:

$$\mathrm{CER} = \sum_{i=1}^{N} E[r_a], \quad \text{where } E\,[r_a] = \begin{cases} 1 & \text{for } \widehat{a} = a \\ E\left[r_a | c_p\right] & \text{for } \widehat{a} \neq a \end{cases}. \tag{7}$$

That is, when the predicted action $\widehat{a}$ is the same as the historic decision a, the reward is 1 (farmers obtain their favorite option). When the predicted action is different, the mean reward conditional on context c is used (farmers do not get their favorite land use, but still receive utility from that option). To compare the results of the models, we normalized the indicator to fall within the range of 0–1.

Second, we included prediction accuracy (i.e., the algorithm predicts farmers/agents' land-use decisions correctly) as an indicator of the model performance. This allowed us to evaluate the algorithm in terms of its ability to reflect farmers' decision, as made in the discrete choice experiment. This is the ratio of the number of correctly predicted instances to the total number of instances.

Third, we analyzed the click-through rate, which measures the accuracy of the recommendations over a subset of farmer-land use pairs that appear in both actual choices and recommendations (Dudik et al., 2011). The click-through rate is important because the goal of our recommender system is not to simply predict farmers' decisions but to pinpoint the recommendations that farmers' would interact with, that is how likely they are to adopt the recommended options. Fourth, the inverse-propensity weighted click-through rate (IPS) assigns weights to land-use options based on the likelihood that they are recommended when the farmer encounters them in the historical data (Dudik et al., 2011). This weighting scheme inverts the probabilities, giving greater weight to less likely actions. Finally, the doubly-robust click-through rate (DR) directly combines predicted values (CTR) with a correction based on the likelihood of an item being recommended by the historical policy when the farmers encounter it in the historical data (IPS) (Dudik et al., 2011). It is supposed to account for potential biases in the data and considers the predicted "click" probabilities based on farmer features (see also supplementary material section S.5).

After computing all five evaluation metrics, we selected the model specification that performed best across these criteria as our land use recommender model for further analysis.

Climate, market, and policy scenarios

We used our trained contextual multi-armed bandit model based on farm-level observed data (compare Sections 4.1 and 4.2) to run several climate, market, and policy scenarios. In our baseline scenario, we adopted the conditions outlined in the discrete choice experiment as the prevailing conditions for each grid cell. A two-step procedure was conducted for the market and policy scenarios. First, we re-trained the model by adding contextual features on top of the local weather history to reduce uncertainty by explicitly accounting for the features (compare Figure 1). Here, we used *contribution margin* and *agri-environmental payments for agroforestry* as additional features (i.e., independent variables) during the training phase. This resulted in two scenarios. In the market scenario, the agroforestry options do not have the potential to earn farmers a higher contribution margin than arable crop farming. Thus, we fixed the profitability of the agroforestry practices at the level of arable farming (€400 per ha) for each grid cell. Likewise, in the policy scenario, we guarantee an agri-environmental payment for agroforestry of €200 per ha (compare Table 1). The two-attribute scenarios were chosen for their comprehensiveness and accessibility, allowing the reinforcement learning approach to be demonstrated effectively. The choice to focus on the contribution margin and agri-environmental payments reflects the importance of market conditions and policy factors in evaluating the adaptability and profitability of agroforestry practices (Augère-Granier, 2020; Gillich et al., 2019). These scenarios provide a solid foundation for exploring more complex scenarios (with more attributes) in the future.

Second, we used the re-trained algorithms to recommend optimal land use under different climate projections. To do so, we supplied the trained recommender models, including the models used for the market and policy scenarios, with climate projection data at the raster level using RCP 4.5 and RCP 8.5 as underlying climate scenarios (Section 5.3). Based on these data, we computed land use recommendations for 2020, 2030, and 2050 under different climate, market, and policy scenarios. Thus, we can assess how changing these features influences the optimal land-use allocation under different climate projections.[11]

TABLE 1 Scenario description.

	Contribution margin (€/ha)		Agri-environmental payment (€/ha)	
Scenario	Agroforestry	Arable crop farming	Agroforestry	Arable crop farming
Baseline	400–800	400	0–200	0
Market	400	400	0–200	0
Policy	400–800	400	200	0

Note: The market and policy scenarios involve incorporating additional contextual variables on top of farms' weather history during training. During prediction, they are kept fixed to reflect specific scenarios at predetermined levels. In the market scenario, contribution margin is included as an additional contextual variable and fixed at €400. In the policy scenario, agri-environmental payment is included as an additional contextual variable and fixed at €200.

RESULTS

We find that introducing agroforestry (here, alley-cropping and short-rotation coppice) can be a viable adaptation strategy for farmers. If introduced in a spatially optimal way, it could increase the overall utility of farmers under given and future climates by more than 20%. In the market scenario (i.e., agroforestry does not earn farmers a higher contribution margin than arable crop farming), we observe a decrease in alley-cropping of approximately 10% compared with the baseline. In contrast, an agri-environmental payment of €200 per ha increases alley-cropping by 1% and short-rotation coppice by 10%. Revealingly, climate change not only implies shifts in the total shares of the land use options but also leads to a spatial re-distribution of the optimal agroforestry land use. In the following section, we describe the selected model before presenting the results of the baseline, and market and policy scenarios (Section 5.2). In Section 5.3, we describe the results of combining policy and climate scenarios.

Allocation model choice

We opted for the decision tree UCB algorithm as our preferred model, with the following specific hyperparameter settings: $\alpha = 0.5$, minimum requirement of two samples to split an internal node, and a maximum depth of 8. When evaluating the performance of various candidate algorithms on the test set, this particular model surpassed all the others. Table S2 and Figure S5 show that the decision-tree UCB algorithm achieved the highest cumulative reward, accuracy, and click-through rate (66%). Additionally, it ranks among the top specifications in terms of the doubly-robust click-through rate.

Optimized land use allocation in space

Figure 3a.1 depicts the optimized land use allocation for the baseline scenario based on the decision tree UCB multi-armed bandit. This implies that under the conditions of the discrete choice experiment (Section 3.1), overall farmers' expected utility could be improved by increasing the share of short-rotation coppice and alley-cropping systems in the agricultural landscape.[12] By comparing this spatially optimized situation with alternative situations, we find that our approach leads to a considerably higher cumulative expected utility given the experimental set-up of the analysis (Figure 3a.2). For instance, an agricultural landscape with only arable crop farming (i.e., without agroforestry) leads to a loss of 19% in cumulative rewards compared with the optimized scenario.

In the market scenario in which there is no prospect of a higher contribution margin, the share of arable crop land increases at the expense of alley-cropping. This is especially true for the fertile regions in the middle of Bavaria (Figure 3b). In the policy scenario, that is, with a guaranteed agri-environmental payment of €200 per ha (Figure 3c), the share of agroforestry goes up to 60% of all available land. These land-use adjustments also affect the cumulative expected reward for farmers, which is the highest in the policy scenario. With the agri-environmental payment, the reward is 10% higher than that of the baseline (see Table 2). However, this would likely result in relatively high costs (compare Section 6.2).

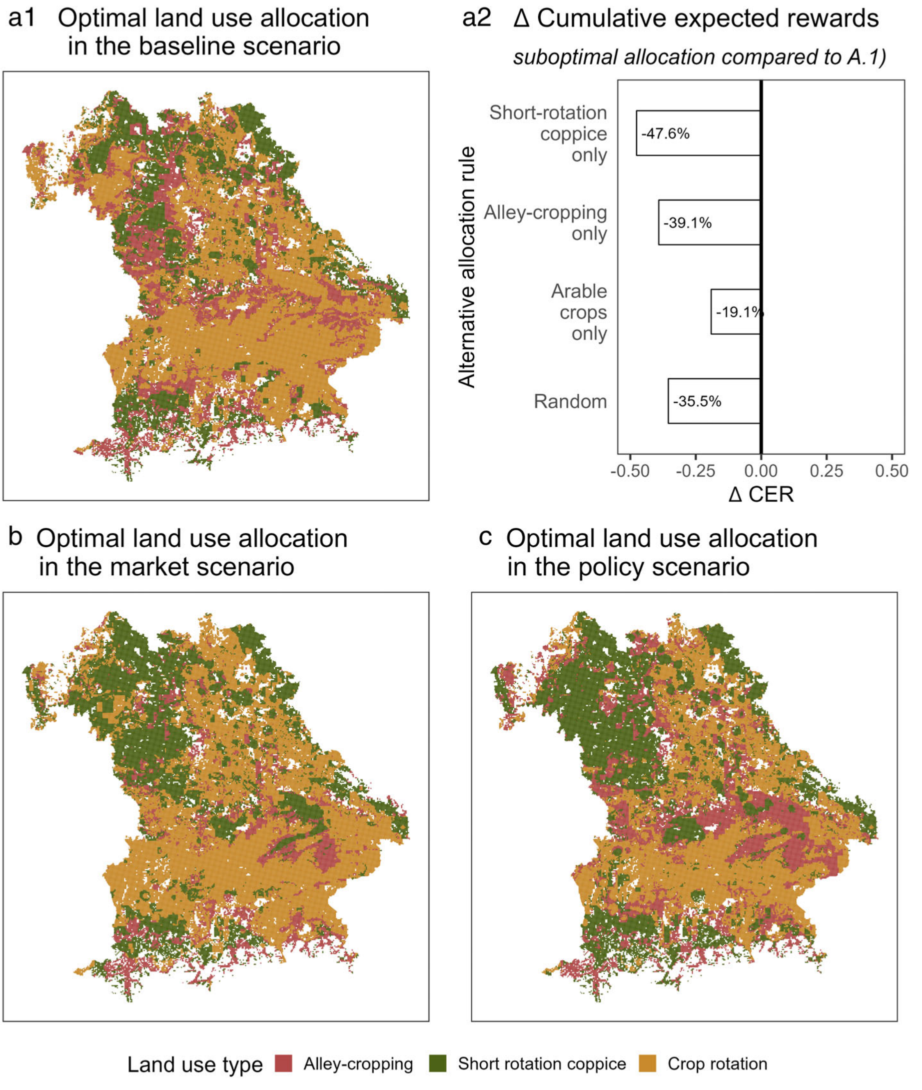

FIGURE 3 (a.1) Optimal land use allocation map for the baseline scenario at current climate based on the decision tree upper confidence bound (UCB) algorithm. (a.2) Cumulative expected rewards of alternative land use allocation rules compared to optimized baseline. (b) Optimal land use allocation map for the market scenario at current climate based on the decision tree UCB algorithm. (c) Optimal land use allocation map for the policy scenario at current climate based on the decision tree UCB algorithm.

Land use allocation under climate projections

Figure 4 shows the potential adjustments in the three scenarios for the two climate projections (RCP 4.5, RCP 8.5). Under the assumed climate projections, we found spatial adjustments over time across projections and scenarios (Figure 4a). Under the RCP 4.5 climate projection and in the policy support scenario, the amount of alley cropping increases under the climate change scenarios from initially less than 12% to approximately 27% in 2050. At the same time, short-rotation coppice declines from more than 56% to less than 37%. Under RCP 8.5, the missing prospect of a high contribution margin for alley-cropping negatively affects its

TABLE 2 Scenario comparison: percentage change in cumulative expected reward (ΔCER) compared to the baseline land-use allocation, and percentage share of land-use types for the different scenarios.

		Share of land-use type (%)		
Scenario	**ΔCER (%)**	**Crop rotation**	**Alley-cropping**	**Short-rotation coppice**
Baseline	0.00	51.89	23.54	24.56
Market	−6.07	53.06	13.36	33.59
Policy	10.03	40.05	24.34	35.61

adoption. Overall, our results suggest that the amount of agroforestry does not necessarily increase under climate change, given farmers' preferences in 2020. However, regardless of the scenario and GHG trajectories, we identified land use shifts over space.

Figure 4c highlights the importance of optimized land use allocation shifts in the face of climate change, as the cumulative expected reward can be maximized. Taking the baseline scenario under the RCP 4.5 trajectory until 2050. By optimizing land-use planning, the expected reward can be increased by approximately 10%. However, maintaining the land use allocation of 2020, the expected reward decreases slightly and is approximately 20% lower than that in the optimized situation. Arable crop farming alone also leads to an increase in the cumulative expected rewards. However, it is considerably lower than that when it is spatially optimized. All the scenarios presented in Figure 4c indicate that spatially optimizing land use leads to higher overall rewards for farmers.

Because climate projections are inherently subject to large uncertainties (Tebaldi & Knutti, 2007), we conducted several sensitivity tests to assess how our results differ across climate projections (see supplementary material section S.6).[13] At the aggregate level, the cumulative expected rewards are, in most cases, within the ±10% range of our model of choice. However, looking at the land-use shares across scenarios, the results vary relatively strongly (Figure S6). For example, if we consider the percentage overlap of raster cells across projections, this metric can be as low as 40% and rarely exceeds 65% (see Tables S3 and S4). In summary, considerable uncertainty surrounds the land-use projections presented in this section and should thus be interpreted with caution. However, general observations such as the relocation of land use in space and the stability of land-use shares over time, still hold true across models.

DISCUSSION

Modeling adaptation to climate change using reinforcement learning

Recommender systems based on reinforcement learning are successfully in use across various industries (Silva et al., 2022). Our findings demonstrate that such a system may offer several advantages also for land use modeling.

First, the mapping of context (here, behavioral and environmental data) to action (here, the adoption of agroforestry practices) provides a novel perspective on the concomitant upscaling of behavioral and environmental information. Upscaling of interrelated natural, institutional, economic, and behavioral factors is an important challenge in land use modeling (Meyfroidt et al., 2018). The algorithm applied here could complement existing modeling approaches at the

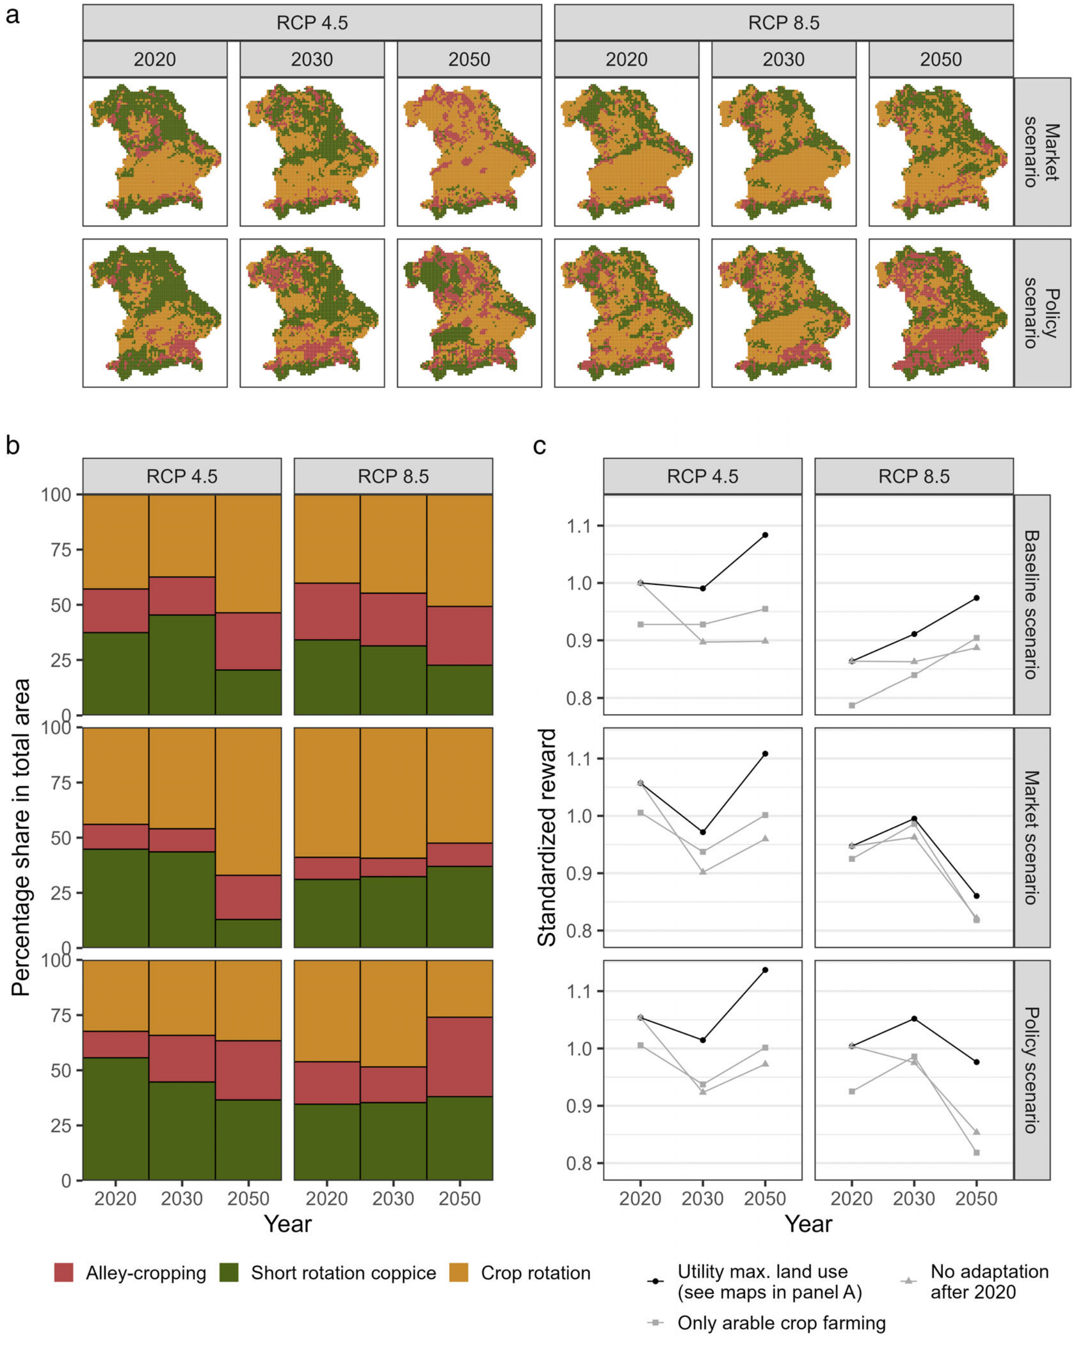

FIGURE 4 Land use allocation under climate projections. (a) Spatial distribution of different land use types. (b) Shares of land use types. (c) Development of the cumulative expected reward given different adaptation assumptions. Rewards are normalized to the baseline scenario RCP 4.5 in 2020. RCP, Representative Concentration Pathway.

macro level by mapping case-study-based local knowledge into larger areas, controlling for the given context. This could provide leverage for in-depth research on the behavioral and economic aspects of land-use decision-making by upscaling information to the spatial level at which policy decisions are taken. The approach could then be understood as an analogy to benefit transfer in environmental valuation (Johnston et al., 2015). By utilizing contextual multi-armed bandits, we can apply data-driven insights to a different context where there is a lack of direct evidence or missing data.

The selected algorithm also allows for computationally efficient integration of behavioral data from an economic experiment with spatially explicit environmental data. The integration of such micro-scale data in a spatially explicit manner remains a major challenge in land-use modeling (Troost et al., 2022). Our proposed approach allows the representation of individual decision-making heterogeneity, capturing a diverse range of preferences that exist at the farm level. This is in line with recent efforts to improve land-use models (Brown et al., 2021; Müller et al., 2020; Müller-Hansen et al., 2017; Troost & Berger, 2014). However, the proposed contextual multi-armed bandits algorithm has significantly lower computational demand. A single model run takes no more than several seconds on a single processor, which is very short compared with, for instance, agent-based models that include individual farm behavior (see Shang et al., 2023; Troost et al., 2022). This offers significant potential for complementing and enhancing spatially explicit impact assessments because it would allow to address the criticism surrounding current aggregate macro-models, which often rely on restrictive behavioral assumptions (Müller-Hansen et al., 2017; Troost et al., 2022).

A significant challenge regarding the application of modern machine learning algorithms is the lack of interpretability of their underlying mechanisms, which makes it difficult to comprehend how they derive their predictions. This can hinder our understanding of the underlying mechanisms driving the predictions of these models, making it difficult to assess their reliability and applicability (Molnar et al., 2020). In recent years, there has been a growing interest in interpretable machine learning, a research field dedicated to developing methods that can explain the decision-making processes of machine learning models (Molnar et al., 2020; Storm et al., 2020). This development presents a promising avenue for exploring new applications of machine learning in agricultural economics, with the potential to enhance decision-making tools and optimize resource allocation.

Next, the land-use simulation approach presented herein is characterized by high flexibility. The flexibility of integrating new data allows for the dynamic inclusion of new actions, such as land-use types, additional explanatory factors, and more observations (Lu et al., 2010). Depending on the research question at hand, land use modelers could also add additional context such as information on farm structures, farmer behavior or environmental data, to train the underlying algorithm. This flexibility is important as it allows the method to be updated in light of the dynamic environment. As the algorithm observes more data and learns from the outcomes, it can dynamically update its decisions to optimize the rewards. This adaptability is beneficial when underlying reward distributions change over time. For example, farmers' preferences could change over time (see Finger et al., 2023). This may not be captured well by alternative methods, such as econometric or supervised machine learning models. Among these models, the learned model is usually static and may not easily adapt to changing conditions without retraining (e.g., Burke & Emerick, 2015; Mérel & Gammans, 2021; Webber et al., 2020).

Finally, this approach provides ample opportunities for adaptation studies. While we focused on land use here, this approach may also be valuable in other settings. This could include modeling the adoption of novel technologies, social innovations, or sustainable farming practices in an uncertain environment.

Recommendations for supporting agroforestry practices in agriculture

Our simulations can provide recommendations for stakeholders (i.e., farmers, advisors, policymakers) to plan climate adaptation using, for example, information, nudges, or policy instruments. Ex ante assessments of climate and policy impacts on agricultural land use show, in a spatially explicit manner, that the adaptation of agroforestry could provide utility to farmers. Advisory services or policymakers could make use of such results by identifying potential hotspots (i.e., areas with a high certainty of benefits) that could be targeted, for example, with spatially explicit advice or by tailoring agri-environmental programs to specific regions.

However, with respect to the policy support scenario, it is important to note that agri-environmental payments might be relatively costly. For example, consider the result of the policy scenario in Section 5.2, which says that farmers' expected rewards could be increased by up to 10% and cultivation area would potentially go up to 65% if farmers received €200 per ha for cultivating agroforestry. Given the more than 200 million ha of cropland in Bavaria, this would potentially result in a rough cost estimate of €223 million per year, which would be approximately 15% of the state budget for food, agriculture, and forestry (Bavarian State Ministry of Finance and Home Affairs, 2023).

The spatially explicit evaluation of agroforestry cultivation potential considering farmers' preferences and local weather history also underscores the importance of spatial planning in addressing climate change impacts (compare e.g., Hurlimann & March, 2012). In our climate scenarios, we found a rather small adaptation using agroforestry (the land use shares would remain approximately constant until 2050). However, the spatial distribution varied significantly in our scenarios. This is in line with Massetti and Mendelsohn (2018) and Cui (2020), who argue that farmers will engage in dynamic spatial adaptation regarding their crop selection as a result of a changing climate; that is, the spatial distribution of crops will change in response to climate change. In summary, our results do not suggest "more" agroforestry in response to climate change but implications for the "where" of agroforestry. Reallocating agroforestry systems can be costly and disruptive. This requires farmers to clear existing crops or livestock pastures, prepare the land for trees, and plant trees (Cardinael et al., 2021). This can be a significant financial investment and can require considerable time and labor (Valdivia et al., 2012). In addition, the reallocation of agroforestry systems can disrupt existing agricultural activities and may require farmers to adjust their production practices. These costs and disruptions can pose obstacles to the ability of farmers to adapt to climate change (Valdivia et al., 2012).

Limitations and further research

The first application of a recommender system in the context of land use modeling inevitably comes with certain limitations. However, as discussed below, the flexibility of the approach provides new entry points for further research.

A critical issue is model performance, as in many other simulation approaches. The accuracy of predicting farmers' decisions was approximately 72%, which was considerably better than that of random classification (54%). Given the fact that our model predicts a complex decision problem, this score is relatively high and similar to those found in comparable tasks such as customer churn prediction, credit risk prediction, medical diagnosis, or crime prediction (Byanjankar et al., 2015; Lalwani et al., 2022; Richens et al., 2020; Safat et al., 2021, e.g.).

Although predictive accuracy is important, it is not always the most important factor for projections because they are often made in complex and uncertain environments. It is difficult to build a model that can accurately predict all the relevant factors, and even if the model is accurate, it may become less accurate over time as the environment changes. We found a click-through rate of 66%, that is, two-thirds of the recommended and consecutively chosen recommendations could be correctly identified in the test data. This value is considerably larger than most click-through rates found in industrial applications (Beel et al., 2016).

However, the performance of the proposed model can be further improved by optimizing the algorithm (Yang & Shami, 2020). This might imply the inclusion of additional data, that is, additional observations and/or additional contextual features that might have an impact on farmers' land use decisions. One promising pathway in this direction is the increasing availability pf open-source data on biophysical characteristics provided at a fine spatial resolution (Chi et al., 2016). As discussed above, this is one of the major benefits of our proposed approach, as it is very flexible in terms of integrating additional data and updating.

Another aspect of our results is that we mapped stated preferences from a choice experiment (rather than observed preferences) to the adaptation of agroforestry. The advantage of using experimental data is that we have a clear and identifiable link between attributes and preferences for the action in our simulation (i.e., the adoption of agroforestry). However, this setting might also imply that we passed through potential biases from the experiment (e.g., selection bias or hypothetical bias) to the simulations. The farmer population may not have the same level of knowledge about agroforestry as the study participants, and may be more pessimistic about potential profitability. This is attributed to the novelty of agroforestry in the study area and the lack of widespread awareness among farmers. In real-world decision-making situations, several factors can limit the adoption of agroforestry, such as within-farm rigidities, perceived profitability, initial negative cash flows, credit constraints, transaction costs, market access, path dependence, and land lease versus ownership (Giannitsopoulos et al., 2020; Glithero et al., 2013; Staton et al., 2022; Valdivia et al., 2012). Further studies could also extend our approach by using observed patterns (e.g., in a quasi-experimental setting) or data from other modeling studies (e.g., more local agent-based models) that would reduce the hypothetical nature of stated preference studies.

While our recommender system effectively utilizes weather/climate conditions and land use attributes to provide recommendations, it currently lacks the ability to incorporate the costs associated with land-use changes owing to the limitations of the cross-sectional data and the absence of an intertemporal component. While we recognize the potential value of this information, obtaining accurate and comprehensive data on land-use change costs is a significant challenge. However, we believe that future developments in the recommender system can be enhanced by incorporating such contextual variables, allowing for the evaluation of the economic costs of different land use decisions. It could also be interesting to explore the potential of combining recommender approaches with investment analysis techniques to gain a more comprehensive understanding of land use decisions and their associated costs.

Another limitation is the model's inability to account for general equilibrium effects. Unlike the recorded historical weather patterns used to train the model, future climate change will not only impact farmers in a specific area and time but also have far-reaching consequences for the entire agricultural sector—nationwide or even globally. These widespread changes inevitably influence agricultural input and output prices, phenomena that a model trained on micro-level data cannot adequately predict. Future research could include general equilibrium effects by developing scenarios that account for such effects.

It is also be important to consider rewards other than farmers' utility. Given the multifunctional nature of farming (Renting et al., 2009), this view may be overly farmer-centered. Nevertheless, our approach allows the adjustment of the reward function to account for example, for ecological sustainability and societal goals, given the required farm- or field-level data (Chi et al., 2016). Finally, another factor that should be considered when interpreting our results is that we had to assume a stable reward distribution until 2050. While it is outside the scope of this study, further scenarios could be evaluated by adjusting farmers' preferences over time.

Another critical issue relates to the chosen case study for the proposed approach. The data are specific to Bavarian farmers, implying a restricted scope of factors and a relatively small sample size. Consequently, the model's direct application to policymaking beyond Bavaria is limited and should be considered when interpreting the empirical results of this study.

Finally, future research should incorporate policy costs associated with land-use changes into their analyses. This can be achieved by gathering more precise predictions of the acreage potentially devoted to each land-use alternative. For example, Gillich et al. (2019) obtained information on the maximum acreage that each farmer was willing to dedicate to novel land-use types through a survey. Modifying the choice experiment to include field-level contextual data would further allow for field-level land use predictions. In both cases, agri-environmental payments could be directly linked to total acreage, potentially providing a more accurate estimate of policy costs.

CONCLUSION

We provide a novel approach to land use modeling based on reinforcement learning. Our proposed method is computationally efficient and spatio-temporally explicit and can account for farmers' behavioral characteristics and biophysical information. To exemplify the potential use of contextual multi-armed bandits in land use simulation modeling, we integrate the findings of a choice experiment regarding novel climate robust land use types, that is, agroforestry, in crop farms located in Southeast Germany. We found that variations in climate, market, and policy conditions could prompt spatial shifts in the potential cultivation of agroforestry systems. Specifically, agroforestry systems have been identified as a viable strategy for climate change adaption. Neglecting to adapt to climate change through adaptive land use adjustments or maintaining conventional arable crop farming practices may lead to significant cumulative utility losses.

Our analysis has important implications for famers, advisory services, and policymakers. First, our results show the potential relevance of agroforestry as adaption strategy to climate change by creating synergies between adaption capacity and farmers' utility. This information can be used to spatially target information campaigns and identify target policy groups. The proposed modeling approach can be used to advance the models currently used for ex ante policy analysis by upscaling existing knowledge on farmers behavioral characteristics and combining it with spatially explicit environmental and farm structural data. This would allow efficiently projecting economic, farm-level land-use decisions into space while considering behavioral factors and assess policy implications. This way, policymakers and stakeholders can gain valuable insights into the spatial and temporal dimensions of decision-making processes. This approach also bridges the gap between micro-level assessments and macro-level projections, enabling them to design context-specific strategies to effectively address climate change adaptation while accounting for regional variations and temporal dynamics.

Our analysis has several implications for future research. The here developed interface of machine learning, experimental economic and land use modeling is an area of increasing relevance. The integration of additional data sources such as information on farm structures, farmer behavior, or environmental data could further enhance the approach in this study. This methodology can be applied to data from other regions, ideally with more variables and observations, to provide valuable insights for local policymakers. Moreover, the application of the proposed method extends beyond land-use modeling, encompass emerging technologies and sustainable farming practices. Finally, integrating the proposed approach with evaluations of agricultural policies holds promise for generating valuable insights into effective policy targeting, thereby driving optimal outcomes in the agricultural sector.

ACKNOWLEDGEMENT

Open access funding provided by Eidgenossische Technische Hochschule Zurich.

ENDNOTES

[1] Tech companies such as Netflix, Spotify, or Google use such recommender systems to provide personalized recommendations. These systems can analyze vast amounts of user data, including viewing or listening history, preferences, or ratings, to generate tailored recommendations for individual users. These recommendations include suggesting movies and TV shows on Netflix, curating personalized music playlists on Spotify, or providing relevant search results on Google (Silva et al., 2022).

[2] It is defined as yearly revenues (i.e., yield times crop price) minus the associated variable cost. It does not include fixed costs or subsidies (compare Gillich et al., 2019).

[3] Reinforcement learning, along with supervised learning and unsupervised learning, forms one of the fundamental paradigms of machine learning (Athey & Imbens, 2019; Storm et al., 2020).

[4] The fact that the context can be described by spatially explicit biophysical information makes it useful for land use modeling. It enables the upscaling of farm-level data into a spatial dimension (compare Section 5).

[5] The order in which the instances are supplied to the algorithm during training does not affect the final outcome since we use all sampled farms to train the algorithm. However, if we were to update the algorithm using novel information, the order would become relevant.

[6] Business-as-usual arable crop farming is defined as standard three-part crop rotation as widely practiced throughout Bavaria and to which farmers can relate well.

[7] Stetter and Sauer (2022) conduct an in-depth analysis of the discrete choice experiment based on the same data as in this study.

[8] In Li and Ortiz-Bobea (2022) a data-driven cross-validation approach was proposed to determine the growing season. This approach is particularly relevant to large regions. However, in the context of the relatively small region under consideration here, the impact of attenuation bias is deemed negligible.

[9] RCPs are trajectories depicting future greenhouse gas concentrations and pollutants resulting from human activities. They quantify greenhouse gas levels and radiative forcing, reflecting the additional energy absorbed by the Earth due to climate change pollution. RCP 2.6 signifies substantial emissions reduction efforts, RCP 4.5 assumes moderate emissions reductions by 2100, and RCP 8.5 represents a trajectory of ongoing rapid greenhouse gas emissions throughout the 21st century (van Vuuren et al., 2011).

[10] Uppercase letters describe matrices.

[11] In this context, it is important to bear the subtle terminological difference between “forecast” and “projection” in mind. Projections are understood as statements about what could happen in the future, given a set of assumptions. Forecasts, on the contrary, are statements about the most likely future outcome, given the current state of the world. Here, we present projections given potential future climate, market, and policy conditions and no anticipation of future events (compare Saaty & Vargas, 1991).

[12] We focus here in particular on weather and climate impacts on farmers' utility. However, there are many more factors that also matter for farmer welfare and decision-making (see also Section 6.3).

[13] It is a challenge to model potential future land-use when there is no historical precedent. This is a difficult problem because it requires us to extrapolate beyond our current knowledge and understanding. However, by using multiple climate projections and accounting thus accounting the uncertainty and variability in climate projections, we can get a better understanding of the range of possible outcomes. However, we emphasize that these projections represent potential scenarios, and not definitive forecasts. Therefore, a literal interpretation of our results should be avoided.

REFERENCES

Antle, J. M., and S. M. Capalbo. 2001. "Econometric-Process Models for Integrated Assessment of Agricultural Production Systems." *American Journal of Agricultural Economics* 83(2): 389–401.

Arneth, A., C. Brown, and M. D. Rounsevell. 2014. "Global Models of Human Decision-Making for Land-Based Mitigation and Adaptation Assessment." *Nature Climate Change* 4(7): 550–57.

Athey, S. 2017. "Beyond Prediction: Using Big Data for Policy Problems." *Science* 355(February): 483–85.

Athey, S., and G. W. Imbens. 2019. "Machine Learning Methods That Economists Should Know About." *Annual Review of Economics* 11(1): 685–725.

Auer, P., N. Cesa-Bianchi, and P. Fischer. 2002. "Finite-Time Analysis of the Multiarmed Bandit Problem." *Machine Learning* 47: 235–256.

Auffhammer, M. 2018. "Quantifying Economic Damages from Climate Change." *Journal of Economic Perspectives* 32(4): 33–52.

Augère-Granier, M.-L.. 2020. "Agroforestry in the European Union."

Bavarian State Ministry for the Environment and Consumer Protection. 2021. "Climate Report Bavaria 2021." https://www.bestellen.bayern.de/application/eshop_app000009?SID=1528315719&ACTIONxSESSxSHOWPIC(BILDxKEY:%27stmuv_klima_012%27,BILDxCLASS:%27Artikel%27,BILDxTYPE:%27PDF%27).

Bavarian State Ministry of Finance and Home Affairs. 2023. "2023 Budget of the Free State of Bavaria." https://www.stmfh.bayern.de/haushalt/2023/haushaltsplan/Gesamthaushalt.pdf.

Bavarian State Office for Statistics. 2023. "Agricultural Census and Agricultural Structure Survey (>2010)."

Beel, J., B. Gipp, S. Langer, and C. Breitinger. 2016. "Research-Paper Recommender Systems: A Literature Survey." *International Journal on Digital Libraries* 17(4): 305–338.

Bouneffouf, D., and I. Rish. 2019. "A Survey on Practical Applications of Multi-Armed and Contextual Bandits."

Brown, C., P. Alexander, S. Holzhauer, and M. D. A. Rounsevell. 2017. "Behavioral Models of Climate Change Adaptation and Mitigation in Land-Based Sectors." *WIREs Climate Change* 8(2): 1–16.

Brown, C., I. Holman, and M. Rounsevell. 2021. "How Modelling Paradigms Affect Simulated Future Land Use Change." *Earth System Dynamics* 12(1): 211–231.

Buhmann, M. D., P. Melville, V. Sindhwani, N. Quadrianto, W. L. Buntine, L. Torgo, X. Zhang, et al. 2011. "Recommender Systems." In *Encyclopedia of Machine Learning* 829–838. Boston, MA: Springer US.

Burke, M., and K. Emerick. 2015. "Adaptation to Climate Change: Evidence from US Agriculture." *American Economic Journal: Economic Policy* 8(3): 106–140.

Byanjankar, A., M. Heikkila, and J. Mezei. 2015. "Predicting Credit Risk in Peer-to-Peer Lending: A Neural Network Approach." In *Proceedings – 2015 IEEE Symposium Series on Computational Intelligence*, SSCI 2015, pp. 719–25.

Cabral, J. S., A. Mendoza-Ponce, A. P. da Silva, J. Oberpriller, A. Mimet, J. Kieslinger, T. Berger, et al. 2023. "The Road to Integrate Climate Change Projections with Regional Land-Use–Biodiversity Models." *People and Nature* 30(1): 106–118.

Cardinael, R., G. Cadisch, M. Gosme, M. Oelbermann, and M. van Noordwijk. 2021. "Climate Change Mitigation and Adaptation in Agriculture: Why Agroforestry Should be Part of the Solution." *Agriculture, Ecosystems & Environment* 319(June): 107555.

Carter, C., X. Cui, D. Ghanem, and P. Mérel. 2018. "Identifying the Economic Impacts of Climate Change on Agriculture." *Annual Review of Resource Economics* 10(June): 361–380.

Chi, M., A. Plaza, J. A. Benediktsson, Z. Sun, J. Shen, and Y. Zhu. 2016. "Big Data for Remote Sensing: Challenges and Opportunities." *Proceedings of the IEEE* 104(11): 2207–19.

Cornes, R. C., G. van der Schrier, E. J. M. van den Besselaar, and P. D. Jones. 2018. "An Ensemble Version of the E-OBS Temperature and Precipitation Data Sets." *Journal of Geophysical Research: Atmospheres* 123(17): 9391–9409.

Cui, X. 2020. "Climate Change and Adaptation in Agriculture: Evidence from US Cropping Patterns." *Journal of Environmental Economics and Management* 101: 102306.

Dobbin, K. K., and R. M. Simon. 2011. "Optimally Splitting Cases for Training and Testing High Dimensional Classifiers." *BMC Medical Genomics* 4(1): 31.

Dudik, M., J. Langford, and L. Li. 2011. "Doubly Robust Policy Evaluation and Learning." In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*, pp. 1097–104.

DWD. 2022. *German Climate Atlas Explanations*. Offenbach: Germany's National Meteorological Service -Deutscher Wetterdienst www.dwd.de/EN/ourservices/germanclimateatlas/explanations/elements/_functions/faqkarussel/heissetage.html.

DWD Climate Data Center (CDC). 2023. "High-Resolution Gridded Daily Data Set (HYRAS)." https://opendata.dwd.de/climate_environment/CDC/grids_germany/daily/hyras_de/.

El Benni, N., C. Grovermann, and R. Finger. 2023. "Towards More Evidence-Based Agricultural and Food Policies." *Q Open* 3(3): 1–24.

ETCCDI. 2018. *Climate Change Indices*. Geneva: Expert Team on Climate Change Detection and Indices http://etccdi.pacificclimate.org/list_27_indices.shtml.

Fei, C. J., and B. A. McCarl. 2023. "The Role and Use of Mathematical Programming in Agricultural, Natural Resource, and Climate Change Analysis." *Annual Review of Resource Economics* 15(1): 1–24.

Fezzi, C., and I. J. Bateman. 2011. "Structural Agricultural Land Use Modeling for Spatial Agro-Environmental Policy Analysis." *American Journal of Agricultural Economics* 93(4): 1168–88.

Finger, R., D. Wüpper, and C. McCallum. 2023. "The (in)Stability of Farmer Risk Preferences." *Journal of Agricultural Economics* 74(1): 155–167.

Fisher, A. C., W. M. Hanemann, M. J. Roberts, and W. Schlenker. 2012. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment." *American Economic Review* 102(7): 3749–60.

Giannitsopoulos, M. L., A. R. Graves, P. J. Burgess, J. Crous-Duran, G. Moreno, F. Herzog, J. H. Palma, S. Kay, and S. García de Jalón. 2020. "Whole System Valuation of Arable, Agroforestry and Tree-Only Systems at Three Case Study Sites in Europe." *Journal of Cleaner Production* 269: 122283.

Gillich, C., M. Narjes, T. Krimly, and C. Lippert. 2019. "Combining Choice Modeling Estimates and Stochastic Simulations to Assess the Potential of New Crops-the Case of Lignocellulosic Perennials in Southwestern Germany." *GCB Bioenergy* 11(1): 289–303.

Glithero, N. J., P. Wilson, and S. J. Ramsden. 2013. "Prospects for Arable Farm Uptake of Short Rotation Coppice Willow and Miscanthus in England." *Applied Energy* 107: 209–218.

Graveline, N., and P. Merel. 2014. "Intensive and Extensive Margin Adjustments to Water Scarcity in France's Cereal Belt." *European Review of Agricultural Economics* 41(5): 707–743.

Hastie, T., R. Tibshirani, and J. Friedman. 2009. *The Elements of Statistical Learning*, Vol 27 Springer Series in Statistics. New York: Springer New York.

Hoerl, A. E., and R. W. Kennard. 1970. "Ridge Regression: Applications to Nonorthogonal Problems." *Technometrics* 12(1): 69–82.

Huber, R., M. Bakker, A. Balmann, T. Berger, M. Bithell, C. Brown, A. Grêt-Regamey, et al. 2018. "Representation of Decision-Making in European Agricultural Agent-Based Models." *Agricultural Systems* 167: 143–160.

Hurlimann, A. C., and A. P. March. 2012. "The Role of Spatial Planning in Adapting to Climate Change." *Wiley Interdisciplinary Reviews: Climate Change* 3(5): 477–488.

IPCC. 2022. *Climate Change 2022: Impacts, Adaptation and Vulnerability*. Summary for Policymakers. Cambridge and New York: Cambridge University Press.

Ji, X., and K. M. Cobourn. 2021. *Weather Fluctuations, Expectation Formation, and Short-Run Behavioral Responses to Climate Change*, Vol 78. Netherlands: Springer.

Johnston, R. J., J. Rolfe, R. S. Rosenberger, and R. Brouwer, eds. 2015. *Benefit Transfer of Environmental and Resource Values*, Vol 14 The Economics of Non-Market Goods and Resources. Netherlands: Springer.

Lalwani, P., M. K. Mishra, J. S. Chadha, and P. Sethi. 2022. "Customer Churn Prediction System: A Machine Learning Approach." *Computing* 104(2): 271–294.

Lancaster, K. J. 1966. "A New Approach to Consumer Theory." *Journal of Political Economy* 74(2): 132–157.

Lehmann, N., R. Finger, T. Klein, P. Calanca, and A. Walter. 2013. "Adapting Crop Management Practices to Climate Change: Modeling Optimal Solutions at the Field Scale." *Agricultural Systems* 117: 55–65.

LfU. 2019. *Bavarian environmental report*. Augsburg: Bavarian Environmental Agency (LfU).

LfU. 2020. "Beobachtungsdaten, Klimaprojektionsensemble und Klimakennwerte für Bayern." 12 pp.

Li, L., W. Chu, J. Langford, and R. E. Schapire. 2010. "A Contextual-Bandit Approach to Personalized News Article Recommendation." In *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, pp. 661–70.

Li, Z., and A. Ortiz-Bobea. 2022. "On the Timing of Relevant Weather Conditions in Agriculture." *Journal of the Agricultural and Applied Economics Association* 1(2): 180–195.

Lobell, D. B. 2014. "Climate Change Adaptation in Crop Production: Beware of Illusions." *Global Food Security* 3(2): 72–76.

Louviere, J. J., D. A. Hensher, J. D. Swait, and W. Adamowicz. 2000. *Stated Choice Methods*. Cambridge: Cambridge University Press.

Lu, T., D. Pál, and M. Pál. 2010. "Contextual Multi-Armed Bandits." *Thirteenth International Conference on Artificial Intelligence and Statistics* 9: 485–492.

Massetti, E., and R. Mendelsohn. 2018. "Measuring Climate Adaptation: Methods and Evidence." *Review of Environmental Economics and Policy* 12(2): 324–341.

McBride, L., C. B. Barrett, C. Browne, L. Hu, Y. Liu, D. S. Matteson, Y. Sun, and J. Wen. 2022. "Predicting Poverty and Malnutrition for Targeting, Mapping, Monitoring, and Early Warning." *Applied Economic Perspectives and Policy* 44(2): 879–892.

McFadden, D.. 1973. "Conditional Logit Analysis of Qualitative Choice Behavior."

McNellis, R., A. N. Elmachtoub, S. Oh, and M. Petrik. 2017. "A Practical Method for Solving Contextual Bandit Problems Using Decision Trees." In *Uncertainty in Artificial Intelligence – Proceedings of the 33rd Conference, UAI 2017*.

Mendelsohn, R., W. D. Nordhaus, and D. Shaw. 1994. "American Economic Association the Impact of Global Warming on Agriculture: A Ricardian Analysis." *The American Economic Review* 84(4): 753–771.

Mérel, P., and M. Gammans. 2021. "Climate Econometrics: Can the Panel Approach Account for Long-Run Adaptation?" *American Journal of Agricultural Economics* 86(5): ajae12200.

Meyfroidt, P., R. Roy Chowdhury, A. de Bremond, E. C. Ellis, K. H. Erb, T. Filatova, R. D. Garrett, et al. 2018. "Middle-Range Theories of Land System Change." *Global Environmental Change* 53: 52–67.

Molnar, C., G. Casalicchio, and B. Bischl. 2020. "Interpretable Machine Learning – A Brief History, State-of-the-Art and Challenges." *Communications in Computer and Information Science* 1323(1): 417–431.

Müller, B., F. Hoffmann, T. Heckelei, C. Müller, T. W. Hertel, J. G. Polhill, M. van Wijk, et al. 2020. "Modelling Food Security: Bridging the Gap between the Micro and the Macro Scale." *Global Environmental Change* 63(May): 102085.

Müller-Hansen, F., M. Schlüter, M. Mäs, J. F. Donges, J. J. Kolb, K. Thonicke, and J. Heitzig. 2017. "Towards Representing Human Behavior and Decision Making in Earth System Models – An Overview of Techniques and Approaches." *Earth System Dynamics* 8(4): 977–1007.

Nelson, G. C., H. Valin, R. D. Sands, P. Havlík, H. Ahammad, D. Deryng, J. Elliott, et al. 2014. "Climate Change Effects on Agriculture: Economic Responses to Biophysical Shocks." *Proceedings of the National Academy of Sciences of the United States of America* 111(9): 3274–79.

Nerlove, M. 1958. "Adaptive Expectations and Cobweb Phenomena." *The Quarterly Journal of Economics* 72(2): 227.

Nolte, C. 2020. "High-Resolution Land Value Maps Reveal Underestimation of Conservation Costs in the United States." *Proceedings of the National Academy of Sciences of the United States of America* 117(47): 29577–83.

Ortiz-Bobea, A., T. R. Ault, C. M. Carrillo, R. G. Chambers, and D. B. Lobell. 2021. "Anthropogenic Climate Change Has Slowed Global Agricultural Productivity Growth." *Nature Climate Change* 11(4): 306–312.

Pathak, S., R. L. Dewangan, and S. N. Mohanty. 2021. "Decision Making under Uncertainty and Problem Solving." In *Decision Making and Problem Solving*, edited by Sachi Nandan Mohanty, 1–12. Cham: Springer International Publishing.

Petersen-Rockney, M. 2022. "Farmers Adapt to Climate Change Irrespective of Stated Belief in Climate Change: A California Case Study." *Climatic Change* 173(3–4): 1–23.

Pielke, R. A. 2005. "Land Use and Climate Change." *Science* 310(5754): 1625–26.

Ramsey, S. M., J. S. Bergtold, and J. L. Heier Stamm. 2021. "Field-Level Land-Use Adaptation to Local Weather Trends." *American Journal of Agricultural Economics* 103(4): 1314–41.

Rauthe, M., H. Steiner, U. Riediger, A. Mazurkiewicz, and A. Gratzki. 2013. "A Central European Precipitation Climatology – Part I: Generation and Validation of a High-Resolution Gridded Daily Data Set (HYRAS)." *Meteorologische Zeitschrift* 22(3): 235–256.

Razafimaharo, C., S. Krähenmann, S. Höpp, M. Rauthe, and T. Deutschländer. 2020. "New High-Resolution Gridded Dataset of Daily Mean, Minimum, and Maximum Temperature and Relative Humidity for Central Europe (HYRAS)." *Theoretical and Applied Climatology* 142(3–4): 1531–53.

Reidsma, P., S. Janssen, J. Jansen, and M. K. van Ittersum. 2018. "On the Development and Use of Farm Models for Policy Impact Assessment in the European Union – A Review." *Agricultural Systems* 159: 111–125.

Renting, H., W. Rossing, J. Groot, J. Van der Ploeg, C. Laurent, D. Perraud, D. Stobbelaar, and M. Van Ittersum. 2009. "Exploring Multifunctional Agriculture. A Review of Conceptual Approaches and Prospects for an Integrative Transitional Framework." *Journal of Environmental Management* 90(SUPPL. 2): S112–S123.

Richens, J. G., C. M. Lee, and S. Johri. 2020. "Improving the Accuracy of Medical Diagnosis with Causal Machine Learning." *Nature Communications* 11(1): 3923.

Rindfuss, R. R., S. J. Walsh, B. L. Turner, J. Fox, and V. Mishra. 2004. "Developing a Science of Land Change: Challenges and Methodological Issues." *Proceedings of the National Academy of Sciences of the United States of America* 101(39): 13976–81.

Rödl, A. 2017. "Short Rotation Coppice: Status and Prospects." In *Encyclopedia of Sustainability Science and Technology*, edited by R. A. Meyers, 71–88. New York: Springer Science+Business Media.

Saaty, T. L., and L. G. Vargas. 1991. "Prediction, Projection and Forecasting: Applications of the Analytic Hierarchy Process in Economics, Finance, Politics, Games and Sports." 254 pp.

Safat, W., S. Asghar, and S. A. Gillani. 2021. "Empirical Analysis for Crime Prediction and Forecasting Using Machine Learning and Deep Learning Techniques." *IEEE Access* 9: 70080–94.

Schmidt, L., M. Odening, J. Schlanstein, and M. Ritter. 2022. "Exploring the Weather-Yield Nexus with Artificial Neural Networks." *Agricultural Systems* 196: 103345.

Shang, L., J. Wang, D. Schäfer, T. Heckelei, J. Gall, F. Appel, and H. Storm. 2023. "Surrogate Modelling of a Detailed Farm-Level Model Using Deep Learning." *Journal of Agricultural Economics* 75: 1–26.

Silva, N., H. Werneck, T. Silva, A. C. Pereira, and L. Rocha. 2022. "Multi-Armed Bandits in Recommendation Systems: A Survey of the State-of-the-Art and Future Directions." *Expert Systems with Applications* 197-(October): 116669.

Speekenbrink, M. 2022. "Chasing Unknown Bandits: Uncertainty Guidance in Learning and Decision Making." *Current Directions in Psychological Science* 31(5): 419–427.

Staton, T., T. D. Breeze, R. J. Walters, J. Smith, and R. D. Girling. 2022. "Productivity, Biodiversity Trade-Offs, and Farm Income in an Agroforestry Versus an Arable System." *Ecological Economics* 191(September): 107214.

Stetter, C., P. Mennig, and J. Sauer. 2022. "Using Machine Learning to Identify Heterogeneous Impacts of Agri-Environment Schemes in the EU: A Case Study." *European Review of Agricultural Economics* 49(4): 723–759.

Stetter, C., and J. Sauer. 2022. Agroforestry Adoption in the Face of Regional Weather Extremes In *96th Annual Conference of the Agricultural Economics Society*. Leuven: Agricultural Economics Society.

Stetter, C., S. Wimmer, and J. Sauer. 2023. "Are Intensive Farms More Emission-Efficient? Evidence from German Dairy Farms." *Journal of Agricultural and Resource Economics* 48(1): 136–157.

Storm, H., K. Baylis, and T. Heckelei. 2020. "Machine Learning in Agricultural and Applied Economics." *European Review of Agricultural Economics* 47(3): 849–892.

Strong, E., B. Kleynhans, and S. Kadıoğlu. 2021. "MABWISER: Parallelizable Contextual Multi-Armed Bandits." *International Journal on Artificial Intelligence Tools* 30(4): 2150021.

Tebaldi, C., and R. Knutti. 2007. "The Use of the Multi-Model Ensemble in Probabilistic Climate Projections." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 365(1857): 2053–75.

Troost, C., and T. Berger. 2014. "Dealing with Uncertainty in Agent-Based Simulation: Farm-Level Modeling of Adaptation to Climate Change in Southwest Germany." *American Journal of Agricultural Economics* 97(3): 833–854.

Troost, C., J. Parussis-Krech, M. Mejaíl, and T. Berger. 2022. "Boosting the Scalability of Farm-Level Models: Efficient Surrogate Modeling of Compositional Simulation Output." *Computational Economics* 62(3): 721–759.

Valdivia, C., C. Barbieri, and M. A. Gold. 2012. "Between Forestry and Farming: Policy and Environmental Implications of the Barriers to Agroforestry Adoption." *Canadian Journal of Agricultural Economics* 60(2): 155–175.

van Noordwijk, M., R. Coe, F. L. Sinclair, E. Luedeling, J. Bayala, C. W. Muthuri, P. Cooper, et al. 2021. "Climate Change Adaptation in and through Agroforestry: Four Decades of Research Initiated by Peter Huxley." *Mitigation and Adaptation Strategies for Global Change* 26(5): 18.

van Vuuren, D. P., J. A. Edmonds, M. Kainuma, K. Riahi, and J. Weyant. 2011. "A Special Issue on the RCPs." *Climatic Change* 109(1–2): 1–4.

Verburg, P. H., P. Alexander, T. Evans, N. R. Magliocca, Z. Malek, M. D. Rounsevell, and J. van Vliet. 2019. "Beyond Land Cover Change: Towards a New Generation of Land Use Models." *Current Opinion in Environmental Sustainability* 38(June): 77–85.

Vermorel, J., and M. Mohri. 2005. Multi-Armed Bandit Algorithms and Empirical Evaluation In *Lecture Notes in Computer Science*, 3720 LNAI, pp. 437–48.

Wang, J., M. Bretz, M. A. A. Dewan, and M. A. Delavar. 2022. "Machine Learning in Modelling Land-Use and Land Cover-Change (LULCC): Current Status, Challenges and Prospects." *Science of the Total Environment* 822: 153559.

Webber, H., G. Lischeid, M. Sommer, R. Finger, C. Nendel, T. Gaiser, and F. Ewert. 2020. "No Perfect Storm for Crop Yield Failure in Germany." *Environmental Research Letters* 15(10): 104012.

Wheeler, S., A. Zuo, and H. Bjornlund. 2013. "Farmers' Climate Change Beliefs and Adaptation Strategies for a Water Scarce Future in Australia." *Global Environmental Change* 23(2): 537–547.

Will, A., and A. Hense. 2008. "The Regional Climate Model COSMO-CLM (CCLM)." *Meteorologische Zeitschrift* 17(4): 347–48.

Wimmer, S., C. Stetter, J. Schmitt, and R. Finger. 2023. "Farm-Level Responses to Weather Trends: A Structural Model." *American Journal of Agricultural Economics* 106: 1241–73.

Wolz, K. J., S. T. Lovell, B. E. Branham, W. C. Eddy, K. Keeley, R. S. Revord, M. M. Wander, W. H. Yang, and E. H. DeLucia. 2018. "Frontiers in Alley Cropping: Transformative Solutions for Temperate Agriculture." *Global Change Biology* 24(3): 883–894.

Yang, L., and A. Shami. 2020. "On Hyperparameter Optimization of Machine Learning Algorithms: Theory and Practice." *Neurocomputing* 415: 295–316.

Zier, C., C. Müller, H. Komischke, A. Steinbauer, and F. Bäse. 2020. "Das Bayerische Klimaprojektionsensemble - Audit und Ensemblebildung." pp. 1–52.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Stetter, Christian, Robert Huber, and Robert Finger. 2024. "Agricultural Land Use Modeling and Climate Change Adaptation: A Reinforcement Learning Approach." *Applied Economic Perspectives and Policy* 1–27. https://doi.org/10.1002/aepp.13448