International Institute for Applied Systems Analysis
IIASA www.iiasa.ac.at

Address: IIASA, Schlossplatz 1, A-2361 Laxenburg, Austria
Email: repository@iiasa.ac.at **Telephone:** +43 (0)2236 807 342

YSSP Report

Young Scientists Summer Program

The Role of Learning in Smallholder Farmers' Decision Outcomes: An Agent-Based Modelling Approach

Cristina I. Apetrei, cristina.apetrei@leuphana.de

Approved by

Supervisors: Nikita Strelkovskii, Nikolay Khabarov, Valeria Javalera Rincón

Program: Advancing Systems Analysis (ASA)

30.09.2021

This report represents the work completed by the author during the IIASA Young Scientists Summer Program (YSSP) with approval from the YSSP supervisors.

It was finished by 30.09.2021 and has not been altered or revised since.

Table of Contents

This research was funded by IIASA and its National Member Organizations in Africa, the Americas, Asia, and Europe.

Abstract

Smallholder farmers' decisions have important implications for the global food system and global environmental changes. Models have been successfully used to understand human decision-making, including that of farmers, and to explore policy interventions towards sustainability. A frequent approach has been to represent economic decisions in aggregate ways and to assume perfect information and utility maximisation. In reality, farmers have to cope with uncertainties about the dynamics of the social-ecological system (SES) that they are part of and they also act in ways that deviate from rational choice theory. As such, learning processes and adaptive behaviour represent an unexplored, but potentially rich avenue for understanding decision-making. In this paper we study the social-ecological outcomes of two such processes that have been suggested as key by sustainability scholars: learning-by-doing and social learning. We expand a pre-existing stylised agent-based model (ABM) of human decision-making within an SES to include learning-by-doing and social learning agents, and we study the impact of their learning strategies on economic, ecological and social outcomes. Our results show that learning agents are able to better match their decisions to the ecological conditions than non-learning agents. In addition, depending on the normative goals pursued, one learning strategy might be more suitable than the other. Lastly, we analyse diffusion dynamics and we find that an initial share of learning-by-doing agents of about 11% might constitute a critical mass for this behaviour to become dominant in a population. This points to research areas of policy relevance that could be explored in future studies.

Acknowledgments

I would like to thank my IIASA supervisors for their efficient and supportive guidance throughout the various phases of this project. In particular, I thank Dr Nikita Strelkovskii, my main supervisor, for his advice in all matters related to the program and for always making time to discuss our project despite his numerous commitments. To Dr Nikolay Khabarov I am grateful for all his pragmatic suggestions, as well as for the energy he brought to our meetings and how he always found something positive to say about my work. I thank Dr Valeria Javalera Rincón for our inspiring discussions and for reminding me that a stylised model does not need to lose connection with its real-world relevance. Although the YSSP 2021 program has taken place mostly online due to the ongoing COVID-19 pandemic, my supervisors' involvement and availability created a warm and collaborative environment that emulated all the advantages that being in person in Laxenburg would have entailed.

Thanks to the activities organised by Dr Piotr Zebrowski, I had the opportunity to meet other colleagues from the ASA program, whose feedback on early presentations of this work was very useful in refining my thinking. I especially thank Dr Matthias Wildemeersch for his rich and constructive comments, as well as literature suggestions.

To my fellow YSSP-ers I am thankful for the collegiality and mutual support. It was a pleasure and honour to meet so many bright and talented researchers from all over the world and it is an invaluable gain that some of them have become friends.

The YSSP organization team has done a wonderful job at keeping our administrative tasks to the minimum and allowing instead for many structured opportunities to learn from and interact with colleagues from the entire institute – a big thank you for that!

I am indebted to colleagues and friends from the Balaton Group for having introduced me to systems thinking and systems analysis, but also for practical advice and inspiration towards pursuing this program at IIASA.

Last, but not least, I thank my Ph.D. supervisors at Leuphana University, Lüneburg, Prof Dr Daniel J. Lang and Prof Dr Henrik von Wehrden, for allowing me the flexibility to pursue this project and my agent-based modelling interests.

Above all, I thank my family and my mother for making it easy for me to focus on my work among daily life responsibilities.

This work has been kindly funded through a IIASA Barry Callebaut fellowship.

About the authors

Cristina Apetrei is Ph.D. Candidate at the Faculty of Sustainability, Leuphana University, Lüneburg, Germany. (Contact: cristina.apetrei@leuphana.de)

1. Introduction

Smallholder farmers' decisions affect all of us, as they relate to the global availability and affordability of food. The International Fund for Agricultural Development estimates that there are 500 million small farms in the world and more than 2 billion people depend on them for their livelihoods (IFAD, 2003). In addition, farmers' decisions have implications on the health of our planet and long-term food security, as they impact the environment in numerous ways, ranging from land-use and land-cover changes, to biodiversity loss, soil nutrient depletion, pressure on the aquatic resources and climate change, among others (Foley et al., 2005; Tilman et al., 2001).

Computational models have long proven to be a useful tool for studying agricultural systems and human-environment interactions more broadly. When trying to understand how farmers shape social-ecological systems, the typical approach has been to take an economic perspective and to model their decisions in aggregate ways and as a direct response to market influences (Brown et al., 2017; Huber et al., 2018; Janssen et al., 2007). This assumes access to perfect information on market conditions, strategy options and the associated payoffs, as well as a rational choice perspective on behaviour. In reality, resource and economic dynamics are often unknown or known partially, either due to inherent uncertainties about underlying processes, such as input availability, environmental variability or price fluctuations, as well as due to social structures and institutions mediating the flow and access to certain information. In addition, even if full certainty about the system dynamics were possible, we know from empirical studies that individual behaviour is rather irrational and sensitive to cognitive shortcuts, experimentation, peer influences, habits, and cultural norms (Simon, 1955; Camerer, 1995; Kahneman et al., 2000). Also, see (Meyfroidt, 2013) for an overview of behavioural theories relevant to understanding social-ecological feedbacks). As such, numerous calls have been made for improving the representation of decision-making within computational models by moving beyond classical rational choice approaches (Parker et al., 2003; Rounsevell et al., 2014; Huber et al., 2018; Schlüter et al., 2012).

The modelling of human behaviour within social-ecological systems beyond classical rational choice is still in its early days. However, significant progress has been recently made in specifying alternative behavioural theories, as well as in designing and parameterising agent decision models using clear theoretical assumptions or empirical data (Groeneveld et al., 2017; Schwarz et al., 2019; Filatova et al., 2013). And yet, one area that remains underdeveloped is the representation of learning processes as part of agents' heterogeneous decision-making.

Learning is closely linked to adaptation and, as such, it has become a normative goal within the social-ecological systems literature (Baird et al., 2014; Armitage et al., 2008). In particular in the case of farmers, adaptation is seen as an important variable characterizing decision-making and it is understood as a set of "adjustments in agricultural systems in response to actual or expected stimuli through changes in practices, processes and structures" (Robert et al., 2016, p. 2). Within this context, learning processes are what drive behaviour as a function of observations from the environment. As Miller et al. (2007) point out, while there is only one way to optimise, there are multiple ways to adapt, and it may be the case that different types of adaptive behaviour are equivalent in terms of outcomes. A key question is, therefore: "can we create a coherent science of adaptive agents?" (Miller et al., 2007, p. 82). To answer this broad question, we posit, learning processes need to be better understood and more often explicitly included in models of decision-making. This proposition is consistent with many reviews calling for more attention to representing learning, risk and uncertainty and social interactions within agent-based models (ABMs) (Huber et al., 2018; Schlüter et al., 2012; Schulze et al., 2017; Bousquet et al., 2004).

In this paper we focus on studying two learning processes that have been often mentioned in the sustainability science and social-ecological systems literatures as playing a key role in human decision-making for environmental resource management: learning-by-doing and social learning. The former stems from adaptive management literature and involves a process of adjusting one's actions based on observed feedback from the environment. The latter is linked to the literature of co-management (i.e. with multiple stakeholders) and it involves changes of understanding rooted in participatory processes and social interactions. We provide a clearer delineation of both concepts in the next section.

Against this background, the main objective of our study is to contribute to the research agenda on understanding adaptive agents in general, and farmer behaviour in particular, in two ways: a) by making a conceptual and methodological contribution as to how learning-by-doing and social learning may be represented in a social-ecological model; b) by advancing our current theoretical understanding of how these processes and their interactions might affect smallholder farmers' decision outcomes.

The first contribution, we hope, is explicit in Section 3 and in Section 4 where we detail our approach to integrate a learning component into an already existing agent-based model of an agro-pastoral system, RAGE (Dressler et al., 2019). The second contribution is realised by asking the following research questions to inform our explorations of the model presented in the "Results" section:

1) How are the outcomes of agents' decisions affected by different learning types?
2) How do different types of learning interact and to what effects?

We situate our efforts at the intersection of several literature streams. Firstly, our approach is grounded in a social-ecological systems perspective, meaning that we understand human-environmental interactions to be more than the sum of social and ecological processes. Instead, social-ecological systems are "integrated systems characterised by strong connections and feedbacks within and between social and ecological components that determine their overall dynamics" (Biggs et al., 2021, p. 5), where relational, co-evolutionary and emergence aspects are key (Preiser et al., 2018; Schlüter et al., 2019).

Secondly, consistently with the social-ecological systems lens (Preiser et al., 2018), the paper takes a "complex adaptive systems perspective" on farming systems. This means, among others, that time is a key variable, there is path dependency and learning is seen as an ongoing and interactive process where decisions at the level of the farm are influenced by the broader context in which it operates (Darnhofer et al., 2010).

Thirdly, in line with the adaptive co-management literature, farmers' learning is not limited to individual experimentation, but can also result from interaction and discussions with others (Munaretto et al., 2012; Darnhofer et al., 2010). Our choice of learning processes to be studied here is thus justified.

A final important observation to be made is that, in our study, we consider both learning-by-doing and social learning as processes taking place at the individual level, represented in the model as one household unit. This is in contrast to learning processes situated at the level of an entire community and expressed as culture, norms and institutions.

2. Theoretical Background

Learning has been conceptualised in many ways and it is beyond our purpose here to provide an extensive review of its understanding in fields as diverse as psychology, economics, organizational management, educational sciences. We will instead focus on the two types of learning that are the subject of our study to briefly situate them within the literatures from which they emerged and to specify, along the lines suggested by Bennett et al. (1992), who learns, what is learnt and to what effect.

2.1. Learning-by-Doing

The concept of learning-by-doing originates from studies of aircraft and ship production in the 1930s when Wright (1936) started plotted produced quantities against costs and observed a 20% reduction of unit costs with a doubling of output (see Dosi et al., 2017). Driven largely by empirical observations, the concept entered economics where similar cost-quantity relationships, i.e. "progress curves", began to be quantified for different products. As one of the assumed driving mechanism behind these cost reductions is that with experience it takes less time to produce one unit of output (Miketa et al., 2004), the notion of "learning curves" gained traction. Other terms that are used interchangeably are "progress curves", "startup curves", or "improvement curves" (Glock et al., 2019). However, Thompson (2011) caution that a "progress curve", in the way originally determined by Wright, is more than a "learning curve" as it allows for other explanations beyond cumulative experience for the observed relation between unit costs and output volume, for instance R&D, product design changes, capital investment.

For learning curves, the relationship between experience indicators and performance indicators is usually expressed by a power function (Dosi et al., 2017). Yet, applications also exist with two-factor learning curves that distinguish between cumulative experience – "learning by doing" and accumulated knowledge – i.e. "learning by searching" (Miketa et al., 2004). Such curves have become mainstream in assessing energy technology policies and in modelling energy transitions. For instance, the concept of "learning by doing", understood as the declining cost of renewable energy substitutes as a function of production capacity, has been used to evaluate the role of backstop technologies and the optimal time for transitioning to alternatives (Jouvet et al., 2012). More recently, "learning-by-searching" has also become increasingly used to refer to incorporating R&D costs in driving energy innovation (Berglund et al., 2006).

In homo-technological systems, the concept of learning-by-doing is also sometimes equated to experiential learning and is related to humans' ability to learn from their mistakes (Bointner et al., 2016). This is a measure of the failure rate of a specific system which is supposed to decline with experience. However, because mistakes can occur also due to other factors than lack of learning, e.g. forgetting, scholars emphasize the need to distinguish between "errors of commission" and "errors of omission" (Bointner et al., 2016).

A closely related concept stemming from psychology is that of reinforcement learning. Its roots go back to Skinner's (1938) operational conditioning and it suggests that negative outcomes will lead to avoiding a specific action in the future, while positive outcomes will make the action reoccur. In contrast to cognitive conscious learning, which allows for beliefs about relationships to be formed, the original understanding of reinforcement learning involves no conscious reflection (Brenner, 2006). However, in practical economic applications, considerations of an automatic response to stimuli have been mostly left aside, and implementations of reinforcement learning have been closer to what Brenner (2006) calls routine-based learning models, i.e. models where "there is a direct connection from the agent's experiences and observations to their behaviour" (Brenner, 2006, p. 908). For instance, a typical way to model reinforcement learning is by assigning higher probability in the future to actions that have proven successful in the past (Arifovic et al., 2004).

In its broader understanding, reinforcement learning is important to our discussion because it opens the door to a broad literature on possible algorithms for updating routines based on experimentation.

Beyond its uses in economics, Thompson (2011) claims that learning-by-doing had already gained a lot of popularity at the beginning of the twentieth century as an educational method, following from the ideas developed by Dewey (1988). Closely related, experiential learning was later coined by Kolb Kolb (1984), the founder of organizational learning, to describe how abstract concepts and generalizations are formed by observation and then tested in new situations (see Miettinen, 2000). Within these interpretations, the concepts of "experience" and "reflection" are central.

Finally, within adaptive management, learning-by-doing and experiential learning have been used interchangeably to refer to knowledge generation processes in systems characterised by uncertainty and environmental change (Lindkvist et al., 2014). As such, learning-by-doing is a structured process of adaptation and it refers to gradual changes in behaviour based on observations of past actions. The subject of learning – who learns – is sometimes left ambiguous, while some authors explicitly claim that learning-by-doing through experimentation manifests itself not only as a change in individual behaviour, but also at the community level (Munaretto et al., 2012). This proposition situates learning-by-doing close to some views of social learning, as explained further below.

For our purposes, we understand learning-by-doing as an individual process of adjusting decisions based on observations from the environment. Heuristics are employed as to how to adjust the decision based on observations. This is first-order learning because the rule by which behaviour is adjusted is not in itself altered. It is also a process that takes place at the level of each individual household, without consideration of external factors or other agents. Lastly, learning-by-doing in this conception does not exclude optimisation, as the goal is still to take a decision that reduces losses. However, it is not optimisation in the sense of estimating future outcomes, but rather as a reactive decision to observations.

2.2. Social Learning

Just like learning-by-doing, the concept of social learning has also been subjected to diverse and often conflicting interpretations.

According to Muro et al. (2008), the conceptual origins of social learning are grounded in the work of Miller et al. (1941) who were the first to propose that individuals observe others and then behave according to formed expectations about benefits and rewards. Later on, Bandura (1977) further developed these insights into his social learning theory emphasizing the role of observing, modelling and imitating others. Whereas earlier behaviourist models assumed that it was not possible to observe factors mediating stimulus and response relations, Bandura's model was a cognitive one, suggesting that an observation would lead to a specific behaviour after some thinking process.

Within economics, imitation has been mostly modelled based on the authors' assumptions and depending on the purposes of the model (Brenner, 2006). Typically, individuals would observe others in their close proximity and employ a certain heuristic on when to imitate and when not. Decision rules could include, for instance, imitating agents with highest performance of those observed or calculating and executing an average behaviour (Brenner, 2006).

Within the domain of agriculture, imitation strategies are recognized as effective ways of minimising risks and they often entail copying decision rules rather than specific farming activities (Le et al., 2012). Particularly in situations of risk or where outcomes are highly uncertain, individuals start considering the experiences of others around them (Nowak et al., 2017). In addition, there is evidence that farmers' decisions are influenced by their social *milieu,* i.e. their networks and interactions with friends and neighbours (Janssen et al., 2007; Hunecke et al., 2017). Although trust relationships and personal networks are important in agricultural decision-making, social mimicry alone might sometimes explain how behaviour spreads (Rebaudo et al., 2011). As innovation diffusion theories suggest, others' behaviours influence decisions when clear benefits of adoption are observed or when there is sufficient adoption of an innovation in a community as to alter perceived norms (Nowak et al., 2017).

These latter insights already suggest a slightly different interpretation of social learning than the one originally advanced by Bandura (1977), where the emphasis is on behaviour occurring and diffusing in social networks rather than as a consequence of individual observation.

Furthermore, within collaborative natural resource management, social learning is often understood as a process taking place during participatory processes (Schusler et al., 2003), and which results in three distinct outcomes: changed knowledge, changed actions and changed actor relations (Beers et al., 2016). Similarly, Reed et al. (2010) provide a much cited definition of social learning as "a change in understanding that goes beyond the individual to become situated within wider social units or communities of practice through social interactions between actors within social networks" (p.6). This perspective on social learning as a process of social change deviates from the conceptualisation of Bandura (1977) as change within an individual (Reed et al., 2010; Apetrei et al., 2021).

A useful way for thinking about these differences has been advanced by Rodela (2011) who identified three perspectives on social learning: an individual-centric perspective refers to changes in personal understanding based on social relations, a network-centric perspective emphasizing changes in practices and relationships at group level and a system-centric perspective describing changes in institutional settings and broader policies. It is especially in relation to the latter two that some authors within adaptive (co)management have linked processes of learning-by-doing to social learning, where experiments are seen as boundary objects facilitating community participation and exchange (Munaretto et al., 2012).

In our study, we will primarily look at social learning as an individual process of imitating other successful agents. However, we also consider a second type of social learning as a diffusion process over the entire network.

2.3. Effects of learning and process interactions

To describe what changes as a consequence of learning processes and the various scales at which learning is situated, the concept of a feedback loop is quite useful.

Le et al. (2012) discuss feedback loops and various types of adaptation in the modelling of land-use decisions in an ABM setting. Building upon the Human-Environment System framework developed by Scholz (2011), they distinguish between a primary and a secondary feedback loop that determine human behaviour by feeding information from the environment. The primary loop refers to how „human agents perceive the status of the environment and react to it", while the secondary loop requires a „reframing of the agent's behavioural program" (Le et al., 2012, p. 84). This echoes conceptualizations of single-loop, double-loop and triple-loop learning. Single-loop learning refers to correcting errors by changing actions based on observed feedback from the environment, double-loop refers to changing existing values and rules that drive actions, while triple-loop learning is about changing the broader institutional and societal context underpinning the set of possible rules/strategies (Armitage et al., 2008; Pahl-Wostl, 2009).

Within adaptive management, some scholars also distinguish between technical learning, aimed at reducing structural uncertainty about the dynamics of the resource, and institutional learning, which is a social process leading to an alteration of the decision architecture (Williams et al., 2016). Technical learning is an adjustment made after monitoring and evaluating the state of the resource, and it is nested within institutional learning. Together, these two processes are also referred to as double-loop learning (Argyris et al., 1978), although note that compared to Le et al.'s (2012) approach, the subject of learning is not necessarily the individual.

While learning loops clarify what the learning can be about, they are not always explicit about who is learning. Diduck (2010) identifies five levels at which learning can take place: individual, action group, organizational, network learning, societal learning. We have also mentioned some of these levels and how they relate to the two core processes that are of interest in our study.

2.4. Linking learning to theories of behaviour

One aspect that still remains unclear is the relationship between various learning processes and theories of behaviour. This is of particular relevance for including learning into agent-based models. The modeller is confronted with four main tasks: finding theories about decision-making, formalizing, implementing and documenting them (Schwarz et al., 2019). Against this background, some learning processes are seen as theories of decision-making in their own right, for instance reinforcement learning and social learning are situated on equal footing with rational choice or bounded rationality meta-theories (Schwarz et al., 2019). More conceptual teasing will be needed in order to advance a framework for how to make systematic choices about incorporating behaviour, and learning in particular, into ABMs. Based on the literature review above, it appears that some learning processes might be compatible with specific behavioural theories, but not with others.

For instance, Schlüter et al. (2017) suggest a list of behavioural theories that can be used as a departure point for modelling human behaviour in ABMs. Among these, the theory of descriptive norms, which assumes that actors will behave in accordance to observed behaviours of others, matches our conceptualisation of social learning at the individual level. Similarly, habitual behaviour theory suggests a response to positive experiences as assumed by learning-by-doing and reinforcement learning. However, at a rather superficial interpretation which is disconnected from the philosophical roots in which those theories emerged, we note that both theories allow for an operationalisation of learning that is grounded in both a rational choice and a bounded rationality view. The theory of descriptive norms tells us that an agent might imitate the behaviour of others, but it does not tell us which others matter in this decision: will the agent attempt to maximise their utility by imitating the most successful agent they see, or will they take a satisficing approach where they will imitate the first agent they meet that performs slightly better, or will they imitate the closest agent in the network regardless of their performance? We hope that the field will advance to offer more tools for disentangling such considerations.

In the case of our study, we implement learning and adaptive agents, but their ultimate motives are still related to utility maximisation. We acknowledge that alternative goals will lead to different implementations, and these will require additional research.

3. Conceptual work

3.1. Towards a framework for modelling farmers' learning-by-doing and social learning

Based on all of the considerations above, we propose here a preliminary and very much simplified conceptual framework for studying farmers' learning-by-doing and social learning in an agent-based model. Below, we briefly reiterate some of the conceptual boundaries we set, but we also highlight operationalisation choices that we made for our model and what alternative options might be considered by modellers in future similar studies.

First, we discuss both learning processes as individual learning, i.e. changes taking place at the level of one agent, in this case one household. Learning-by-doing is a change of actions in response to observed feedback from the environment and a cognitive approach to information processing is assumed. We distinguish among several types of social learning. *Social learning 1* is also modelled as a change at the level of a household, but the trigger is an observation of other agents' outcomes.

Social learning 2 refers to a double-loop type of learning, where agents learn by imitation and alter their decision-making algorithm (i.e. they adopt a different learning type / behavioural strategy). The effects of this learning about learning are manifested both at the individual level as a change in behavioural strategy, but also at the community level because the initial distribution of learning types in a heterogeneous population changes as a consequence of strategy switching. The effects at the community level have an emergent character and they depend upon the changes happening at the individual level. As we will show further, in our analyses of Social learning 2 we will only evaluate the impacts of community level effects (Table 1).

Table 1. Conceptual framework for modelling learning-by-doing and social learning. Highlighted rows indicate the model analyses that we address in this study.

Model elements	Learning types	Level of learning effect	Target of learning (what changes)	Explanatory process description
Behavioural strategy 1	**Learning-by-doing**	Individual level	Agent actions	Cognitive processing of information from the environment
Behavioural strategy 2	**Social learning 1**	Individual level	Agent attributes	Imitation of other agents' characteristics / heuristics
Strategy switching	**Social learning 2**	Individual level	Agent learning type (i.e. behavioural strategy)	Imitation of other agents' learning type
		Community level	Distribution of learning types in the simulated world	Diffusion of learning types

Second, we note that, in our implementation, the drivers behind the explanatory processes described in Table 1 are consistent with an economic view where agents pursue the goal of profit maximisation. In future models, alternative goals could be used as driving the learning behaviours presented.

Third, although this issue is not explicit in our theoretical considerations of learning, agent heterogeneity in dealing with uncertainties is an important aspect that needs to somehow be captured when modelling learning behaviours. For instance, pest dynamics models revealed that there are broad differences in how information is diffused, perceived and used (Rebaudo et al., 2011). In our model we include such heterogeneity by introducing a stochastic "resistance" parameter to mediate behavioural responses to feedback (see Section 4.2).

3.2. Eliciting the relationship between uncertainty and learning

As it has already emerged from our literature review above, learning in environmental management is a process aimed at reducing uncertainty. Within the context of adaptive management, Williams et al. (2016) discuss four sources of uncertainty: *structural uncertainty*, *environmental variation, partial* *control* and *partial observability*. *Structural uncertainty* refers to a limited or no understanding of the underlying dynamics that governs how the resource state changes from one time step to the next. *Environmental variation* includes external factors that affect resource dynamics, for instance precipitation patterns. *Partial observability of the resource* may be linked to problems of access, but we also add here distorted or noisy information flows, as well as what other authors call epistemic uncertainty, i.e. limited knowledge of the state of the resource due to improper measurement or insufficient data, among others (Regan et al., 2002). Finally, *partial controllability* denotes a difference between the intended effects and those that actually occur. This can be due to properties of the agent (attitudes, limited cognitive abilities, errors), but also to other factors affecting the resource state, for instance when multiple users manage the same resource.

Within the context of agricultural systems and our discussions of two learning processes, we add to the typology above another source of uncertainty, i.e. *partial observability of the social conditions*. This is particularly relevant in the context of farmers' imitation behaviour, as certain aspects behind others' performance or decisions might not be fully accessible, behaviour might be difficult to copy (Le et al., 2012) or strategies might be difficult to infer from observations (Miller et al., 2007). Furthermore, the observability of the broader social environment depends on the structure of one's personal network.

In Figure 1, we elaborate on Williams et al.'s (2016) work to represent these sources of uncertainty and their relationship to learning-by-doing and social learning. We argue that these considerations of uncertainty are playing an important role in the micro-decisions made by the modeller when implementing learning processes. For instance, assumptions need to be made about the extent to which an agent can correctly sense the state of an environmental resource. As such, a transparent acknowledgment of these uncertainty sources and how they are handled is necessary.

Due to emergence properties at the level of the coupled social-ecological system, other sources of uncertainty might be relevant (see Schlüter et al., 2019), for instance not knowing whether the behavioural strategy observed socially might be correct when applied to one's situation. For simplicity, we do not explicitly address these in our scheme, but these emergent sources of uncertainty should also be kept in mind and reported as relevant.

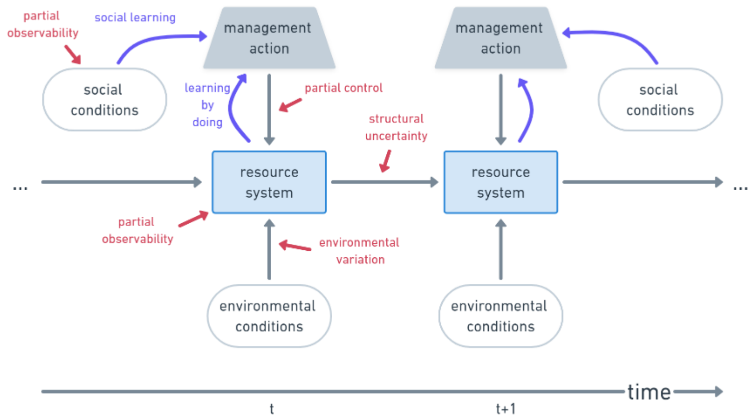

Figure 1. Sources of uncertainty in adaptive management and their relationship to learning. Adapted from Williams et al. (2016).

4. Methods

To realize our goal of advancing current understanding on how learning-by-doing and social learning affect smallholder farmers' decision outcomes, we implement and study these processes in a pre-existing agent-based model of agro-pastoralist communities, RAGE. Agent-based modelling is particular suited to our purposes because it allows for implementing individual heterogeneity as discussed in Section 3.1., but also because it permits the observation of emergent behaviours, thus allowing us to study the effects of social learning 2 at the community level.

We now proceed with a short description of the original RAGE model, followed by details on our learning extension.

4.1. Original Model Description

RAGE is a rangeland grazing model by Dressler et al. (2019), which has been explicitly developed to study diverse theories of human decision-making and their impacts within the context of resource management. Although inspired by empirical work, it is a **stylised** social-ecological model of resource use in a rangeland system, where both the ecological and the social dynamics are kept as simple as possible. The model was developed in NetLogo and it is made available Open Source, via the ComSES database: https://www.comses.net/codebases/5721/releases/1.0.6/.

Our choice of using this model to address our research questions is justified as following. First, we were looking for a stylised model that allowed us to focus on the implementation of learning processes rather than on minute parameters describing very complex resource use dynamics. The very nature of the learning processes, and in particular of learning-by-doing, required that the model is a social-ecological one and that some feedback from the environment does exist. However, a too complicated model would have made it difficult to tease out impacts of learning from impacts of other variables. Second, the rangeland system was easy to repurpose to our thematic interest in smallholder livestock farmers. Third, because the model has been previously developed and used to study behavioural strategies beyond rational choice, it allows for future extensions of our work along the lines we suggested in Sections 2.4 and 3.1. Fourth, the model is a spatial one, which allows us to implement social learning based on observation of neighbours – a diffusion process which is typical for agricultural communities. Fifth, the RAGE model already includes a component of collective action, as it allows for certain social norms to be activated (e.g. pasture resting). Collective action and institutional emergence are aspects closely linked to social learning and they can also make the object of follow-up research to our work here (see Section 6). Lastly, the sharing and reuse of ABMs is encouraged within the SES literature as a way to enhance verification and transferability of insights (Schulze et al., 2017).

In short, the RAGE model comprises of a social component and an ecological component that interact through various feedbacks. The agents are individual households who own livestock that are placed on pastures to graze. The pasture provides fodder for the livestock, and the livestock's grazing affects, in turn, the amount of biomass on the pasture. The regeneration of the pasture is driven by a simple vegetation regrowth model which also depends on precipitation. There are two types of vegetation driving the pasture regeneration dynamics: green biomass and reserve biomass. Green biomass represents those parts of the plants that are easier to regenerate; these are consumed first. Reserve biomass represents the stock of stems and roots of the plants which, when affected, significantly slow down the regeneration of green biomass. The model is run over several time steps. At each time step, households sense the available biomass on surrounding pastures, as well as other information about resting state of certain pastures and they make decisions about where to place their livestock, within a knowledge radius, in such a way that they achieve the goals specific to their behavioural type. Three behavioural types are implemented, each with its own theoretically-informed assumptions: traditionalist, profit maximiser and satisficer. Every year, livestock numbers increase through reproduction. Livestock heads exceeding the fodder availability on the pasture where they are placed die. When a household reaches livestock 0 it is removed from the world. Full details about the model components, including scheduling and the vegetation submodel, can be found in the ODD+D protocol included with the model at ComSES, also available as supplementary information for Dressler et al. (2019).

4.2. Model Extension and Design Choices

We implement learning processes within RAGE by designing an extension module that can be turned on and off from the visual interface in NetLogo. We explain the main aspects of our extension below.

Model overview

Figure 2 illustrates several modifications that we make to the main model components and the overall dynamics of the system.

A first and important design choice that we made was to transform the model from rangeland to an agricultural system which allows individual households to learn solely from the environment and independently of the actions of other agents (unless this is done explicitly through social learning). As such, the decision that agents have to make each round is no longer about *where* to place their livestock, but rather on *how many livestock* to place on their own pasture. Each household exploits the pasture (patch) on which it is situated. The property regime thus also changes relatively to the original model from a common pool resource to a private good. Up to 100 households are randomly distributed across a 10x10 world at the beginning of the simulation. They are also endowed with the same number of initial livestock. The number of initial households and the number of initial livestock of each household are input parameters.

Because all pastures are initialized with the same amount of green and reserve biomass and are all governed by the same resource dynamics, this design decision permits comparisons among agents' performances as a result of their learning. Furthermore, stochasticity associated with precipitation is eliminated. Rain can be easily reactivated in the code, if needed, but learning effects under *environmental variability* (as discussed in Section 3.2.) do not make the object of our study.

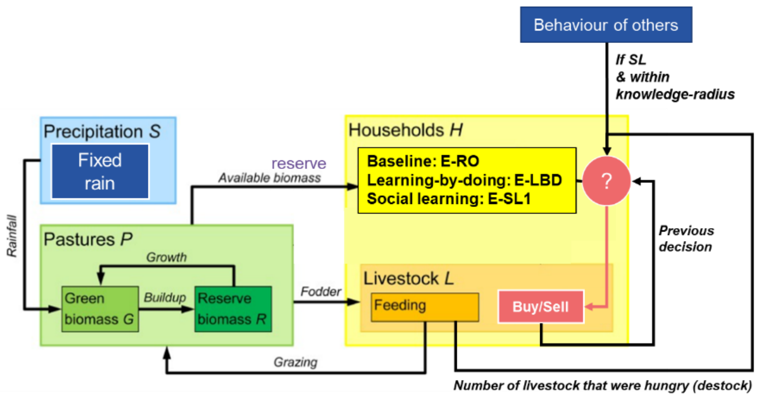

Figure 2. Schematic representation of the modified RAGE model with the learning extension activated. Adapted from Dressler et al. (2019).

Decisions on *how many livestock to place* on the pasture are not limited numerically, as it is assumed that households have full control over the size of their herd and they are able to buy and sell as much livestock as they need to achieve the desired number. Consequently, reproduction of livestock is no longer a relevant variable driving the size of the herd.

Like in the original model, livestock placed on the pasture is feeding on biomass. First they consume the green biomass; if this is over, then they start consuming from the reserve biomass up to a percentage *gr2*, representing the grazing pressure. This gr2 parameter is fixed at 0.1 in our extension, meaning that the maximum reserve biomass that can be consumed at any time is 10% of the remaining available reserve biomass. The amount of fodder needed by each animal is fixed and the same for all households and across all time steps. If in a round more fodder is needed than the sum of available green biomass and 10% of the remaining reserve biomass, the exceeding amount is used to calculate how many livestock went hungry (this information is stored in a variable from the original model called "destock" and, under certain circumstances, it influences household decisions in the next round).

The decisions of how many livestock to place on the pasture each round are affected by: a) the sensed available reserve biomass; b) the household's behavioural strategy (agent learning type); c) previous decision on how much livestock to buy/sell; d) the number of livestock that were observed to be hungry in the previous round (destock value); and e) for agents for which social learning is applicable: the behaviour of neighbours who live within a knowledge radius *k*. The parameter *k* defines the size of the neighbourhood within which social information is searched for under social learning settings. In our model runs we use a value of *k*=1 as default, which defines the Moore neighbourhood of the agent. The *k* parameter stands as one way to model the *partial observability of the social conditions*.

There are two sources of stochasticity in the extension model: 1) the distribution of the r-parameter in the population and 2) the spatial distribution of agents on the map. The scheduling of the different model processes is as following: at the beginning of each time step observations are made (about pasture condition or other agents' state), then household decisions are made, livestock is placed on the pasture and feeds, the number of livestock that went hungry is calculated, the vegetation submodel updates the state of the pasture to be sensed in the next time step.

Agent memory

An important variable in the decision-making of learning agents is information from current and past observations, as well as about past decisions. That is why, the extension adds for all agents a memory of observations. Since the learning algorithms that we implement here are incremental, memory length is 1, i.e. only observations from time step *t-1* are used. Specifically, information is stored about: how many livestock were hungry at the end of the last time step, how many livestock were placed last time (previous decision), what was the amount of observed reserve biomass at the beginning of the last time step.

Agent heterogeneity in dealing with uncertainty: the r-parameter

As we briefly touched upon in Section 3.1, we quickly realised that modelling learning required us to make some explicit choices about agent heterogeneity in preferences as well as in dealing with various types of uncertainty.

Many scholars have drawn attention to such considerations. For instance, Huber et al. (2018) emphasize that farmers' heterogeneous decision-making should capture not only cognitive processes or social interactions, but also the socio-economic and natural context in which they take place, such as opportunity costs for non-agricultural activities and various short-term and long-term calculations. Similarly, Darnhofer et al. (2010) highlight that farmers' choices are constrained by their personal characteristics and external structures, which makes learning a relational understanding of reality rather than an objective cognitive process. Other scholars call for attention to the role of risk attitudes in relation to learning (Marra et al., 2003) or to inaccuracies in payback calculations (Muelder et al., 2018), the latter pointing to our earlier notion of *partial controllability*. Brenner (2006) warns that neglecting individual differences may entail important implications on model outcomes.

But how to organize and operationalise such a broad spectrum of variables that may result in individual differences? A solution we found was to introduce a so-called "resistance-to-learning" factor to capture variability in how agents respond to the information that they acquire from their social-ecological environment. This resistance is different for each agent and it represents a deviation from what the "perfect" learning effect of a specific learning algorithm would be.

In their generic ABM of social learning, Nowak et al. (2017) use a similar mediating variable between observed outcomes and behaviour which they call "propensity to engage in a particular behaviour". In their case, propensity is modelled as a probability of making a particular decision and it "incorporates beliefs, norms, self-efficacy, and intention, as well as other external factors such as access and barriers" (p.5). Individual propensity values can also be influenced by social processes.

In contrast, we model resistance as a percentage r by which agents would deviate from the "rationalized", learned value of how much livestock to place on the livestock. The *r parameter* can take any value between -0.95 and +0.95 in increments of 0.05 and it is randomly drawn from a discrete uniform distribution and fixed for each agent at the beginning of the simulation. It can represent either uncertainty in the sensing (*partial observability of the resource system*), noise in the information received, or an inherent characteristic of the agent, such as a risk attitude. In a world of perfect information and perfect action, the "rationalized" learned decision would be the one for $r=0$. In reality, and depending on the value of their *r-parameter*, agents might over-compensate or under-compensate relatively to the rationalized number calculated via their learning algorithm, by deciding to place more or less livestock, respectively, by a deviation of up to 95%. The choice to limit the range of the *r-parameter* to -0.95 and +0.95 instead of -1 and +1 was made in order to avoid a situation where an agent deviates from the "rationalized" decision by giving up the entire livestock herd and essentially removing himself from the farming activity.

Agent behavioural types

The extension model does not make use of the original behavioural types, but instead includes three new types of agents: baseline, learning-by-doing, social learning 1.

Firstly, **a baseline agent type** is defined, which takes decisions without any learning (extension-r-only, **E-RO**). This type of agent essentially follows the routines of the maximising strategy from the original model (MAX), but the main difference is that an *r-parameter* is included as attribute and consequently E-RO agents are heterogeneous in their actions. This design choice was necessary to enable comparability of outcomes with the learning-by-doing and social learning agents.

The behaviour of baseline agents is represented in Figure 3, where t represents the current time step, $L(t-1)$ is the number of livestock placed in the previous round, $L_d(t-1)$ is the number of livestock that were observed to be hungry in the previous round (i.e. overshooting the carrying capacity of the pasture), and L_R is the "rationalised" number of livestock to be placed, i.e. the value to be placed when $r=0$. L_d is a variable that the agents can sense directly and its value from the previous round is stored in their memory. L_R is calculated based on the routine from the original model of simply destocking livestock exceeding the capacity of the pasture, i.e.:

$$L_R = L(t-1) - L_d(t-1) \quad (1)$$

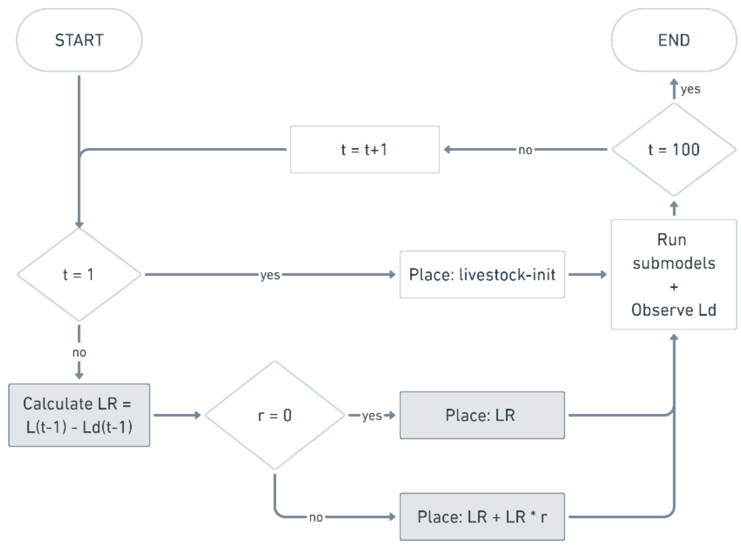

Figure 3. Schematic representation of E-RO agent behaviour, i.e. baseline / no learning agents

When $r \neq 0$, the number of livestock to be placed on the pasture, L, is a deviation by r from L_R:

$$L = L_R \times (1 + r) \tag{2}$$

Secondly, we define **learning-by-doing agents (E-LBD)**, i.e. agents with behavioural strategy 1 as shown in Table 1. Learning-by-doing agents employ a more sophisticated algorithm for deciding how much livestock to place on the pasture which is based on observing the differential changes in the amount of reserve biomass available on the pasture and responding to those with a proportional change in their herd size. Thus, if based on their observations they calculate that the amount of reserve biomass has declined between the previous and the current round by 6%, they will decide to also reduce their herd size by 6% (by selling 6%), with certain further specifications to account for the previously observed hungry livestock. If, however, the observation is that the amount of reserve biomass has increased, then the herd size will also increase proportionally. Figure 4 illustrates the details of the E-LBD behaviour.

The first thing that a learning-by-doing agent is doing is to calculate how many livestock have to be bought or sold as a response to the observed change in reserve biomass, according to the following formula:

$$L_{sb} = \frac{R(t) - R(t-1)}{R(t-1)} \times L(t-1), \tag{3}$$

where $R(t)$ is the reserve biomass observed at the beginning of the current time step, $R(t-1)$ is the reserve biomass observed at the beginning of the previous time step, and $L(t-1)$ is the decision taken in the previous time step. Then, if no overshoot of the carrying capacity has been observed (i.e. L_d=0, no livestock went hungry in the last round), the "rationalized" decision for the new size of the herd L_R is to increase or decrease the previous herd size by L_{sb}. However, if an overshoot of the carrying capacity has been observed ($L_d \neq 0$), then the "rationalized" decision will be to reduce the herd size at least by the observed L_d value. This is represented as following:

$$L_R = \begin{cases} L(t-1) + L_{sb}, & L_d(t-1) = 0 \\ L(t-1) + \min[L_{sb}, -L_d(t-1)], & L_d(t-1) \neq 0 \end{cases} \tag{4}$$

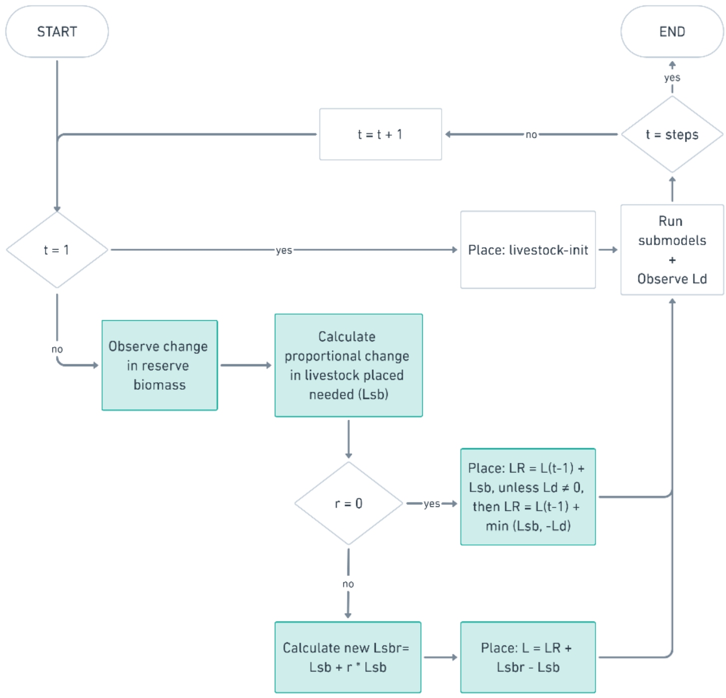

Figure 4. Schematic representation of E-LBD agent behaviour

For instance, even if the change in the state of the pasture might correspond to a decrease of livestock, L_{sb}, of 5 headcounts, but 6 headcounts were observed to be hungry (L_d), then the new "rationalised" decision will be to reduce the previous herd size by 6 and not by 5 headcounts. This accounting of the destock value that we introduce aims to reconcile a rather mechanistic cognitive algorithm of responding to observations in the environment with the conscious planning that might be realistically expected from a farmer.

Further, just like with baseline agent types, the heterogeneity introduced by the r-parameter means that agents will deviate from L_R. We explain how this is implemented, i.e. what happens when $\boldsymbol{r \neq 0}$.

First, a deviation from the "rationalized" number of livestock that should be sold or bought (L_{sb}) is calculated, in a way that the sign (the direction of change, buying vs. selling) is preserved, as following:

$$L_{sbr} = \begin{cases} L_{sb} \times (1 + r), & L_{sb} < 0 \\ L_{sb} \times (1 - r), & L_{sb} \geq 0 \end{cases} \tag{5}$$

Then, the final decision of how much livestock to place on the pasture, L, is a function of the "rationalized" value, L_R, L_{sb} and L_{sbr}. When $r=0$, the final decision is the "rationalized" decision L_R, as shown below:

$$L = \begin{cases} L_R, & r = 0 \\ L_R + L_{sbr} - L_{sb}, & r \neq 0 \end{cases} \tag{6}$$

Thirdly, we model **social learning agents (E-SL1)** to represent "Social Learning 1" from our conceptual framework (Table 1). These are agents who compare their own performance to the performance of their neighbours (from the neighbourhood defined by the knowledge radius k). If there are neighbours who are performing better, then the agents will imitate the *r-parameter* of the most successful neighbour, i.e. the one with the highest number of livestock. If there is no neighbour who performs better, then the agent behaves according to baseline behaviour E-RO. The algorithm is shown in Figure 5.

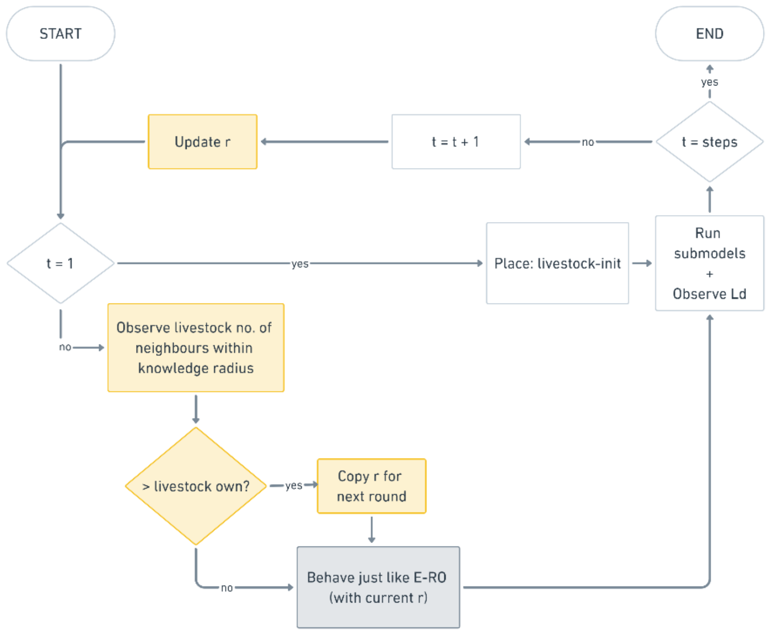

Figure 5. Schematic representation of E-SL1 agent behaviour

The rationale for our design choice to let agents copy their neighbours' r-parameter as opposed to, e.g. their decisions, is that the r-parameter is what drives differences in success. Consider what happens when following Figure 6 given that all agents are initialized in our model with the same pasture conditions and the same initial herd size. At time step 1, all agents will place their entire herd on the pasture (*livestock-init*). Then, at time step 2, all agents will have performed the same, so they will all continue with the baseline behaviour. It is only at time step 3 that social comparison will result in some differences providing scope for imitating others' behaviour. But by then, the pasture conditions will also be different for each agent, so that simply copying others' decision of how much livestock to place on the pasture would not provide the same results as for the agent that is being copied. In addition, as we are trying to model a social learning process, we assume that there would be some form of information exchange about the underlying mechanism of the decision and that it is that mechanism that would be copied, rather than the outcome of the decision itself (the exact number of the livestock to be placed). This is also consistent with the observation by Le et al. (2012) discussed previously that farmers more easily copy decision rules rather than specific activities.

Strategy switching as "Social Learning 2"

Implementing two types of learning agents and a baseline behaviour allows us to compare the social-ecological outcomes of one learning strategy over the other, so as to address our first research question. In relation to our second research question, we also implemented the possibility for a global behaviour for all agents: strategy switching. When this option is activated, all agents have the possibility to select a different learning algorithm, i.e. to change their agent behavioural types, by copying the more successful neighbours. This introduces the possibility to study a second-order type of social learning, which we previously called "Social Learning 2", where the very cognitive rules underlying behaviour are changed. A representation of strategy switching is shown in Figure 6.

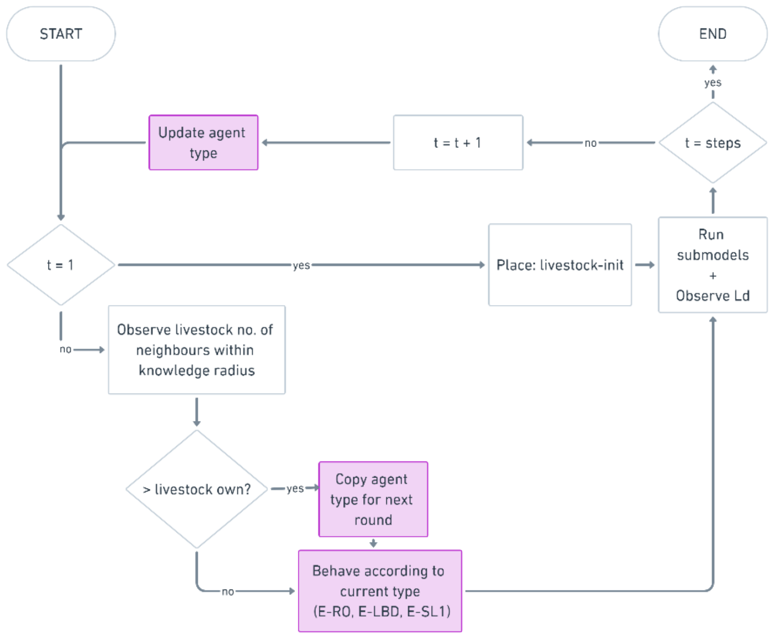

Figure 6. Schematic representation of Social Learning 2 behaviour: strategy switching

The effects of strategy switching can be studied in terms of individual performance compared to performance of agents with no strategy switching, but the more interesting effect that we focus on in our analysis pertains the community level (as highlighted in Table 1), specifically the diffusion of learning strategies (agent types) in a mixed population of E-RO, E-LBD and E-SL1 agents.

As it is now implemented, strategy switching happens whenever the agents observe a neighbour performing better in terms of livestock, using the same criteria as E-SL1 agents who copy the *r-parameter*. As an alternative for future models, social learning 1 and social learning 2 could also be modelled by adding a probability parameter and a logistic function to determine the frequency at which behaviour might change (see e.g. Nhim, 2018). For our purposes, we decided to start with the simplest possible behaviour.

4.3. Experiments and Model Settings

To answer our research questions we explore our model using a design-of-experiments (DOE) approach, as described by Lorscheid et al. (2012). DOE is a systematic process for planning and conducting model runs so that reliable conclusions can be drawn about the relationships between input parameters, model outputs and the processes behind.

In a first step, we defined three sets of experimental objectives corresponding to our refined research questions and we classified our variables of interest (

Table 2

Table 2). In Table 2, "homogeneous agent types" means that the model world is initialised with all agents of the same type, while "heterogeneous agent types" means that it is initialised with a mixed population of all three agent types.

Table 2. Variable classification for three main experiments and the corresponding refined research questions

	Independent variables	**Control**	**Dependent variables**
Experiment 1. **(RQ1)** **Homogeneous agent types**	Type of learning (E-RO vs. E-LBD vs. E-SL1)	Population size Initial herd size For SL: knowledge radius Model seed r-parameter Repetitions Time steps	Economic outcomes: mean total livestock healthy Ecological outcomes: mean reserve biomass Social outcomes: inequality (Gini-index)
Experiment 2. **(RQ2a)** **Homogeneous vs. heterogeneous agent types**	Degree of homogeneity of population (E-HET vs. E-HOM)	Initial distribution of learning types (counts)	Economic outcomes: total livestock healthy Ecological outcomes: mean reserve biomass Social outcomes: inequality (Gini-index)
Experiment 3. **(RQ2b)** **Heterogeneous agent types**	Interaction of learning types (strategy switching, SL2)	Initial distribution of learning types (counts) Population size	Final distribution of learning types (agent type counts)

Research Questions

(RQ1): How are the outcomes of agents' decisions affected by different learning types?

(RQ2): How do different types of learning interact and to what effects?

a) How do agents' decision outcomes differ in a population with homogeneous learning types vs. a population with heterogeneous learning types
b) How does learning-by-doing diffuse in a population of mixed learning strategies?

Experiment 1

For our first research question (RQ1), "How are the outcomes of agents' decisions affected by different learning types?", we are interested to compare the social-ecological outcomes of agents employing different types of learning.

For each of three different experimental treatments, we look at three types of outcomes – economic, ecological and social – operationalised as following: total livestock healthy over the entire period of simulation, mean reserve biomass at the end of the simulation and the Gini-index over the entire population. The *total livestock healthy* of agent *i* (T_i) variable captures the total number of livestock that an agent managed to sustain on the pasture throughout all the time steps of the simulation. Thus, it is the sum of the livestock sustained at each time step, the latter given by the difference between the livestock placed (L) and the livestock that went hungry (destock, L_d):

$$T_i = \sum_{t=1}^{n} L_i(t) - L_{d_i}(t), \tag{7}$$

where *i* is the agent and *n* is the total number of time steps. We quantify economic outcomes at the level of the entire population of agents. As such, our dependent variable of interest is the *mean total livestock healthy*, μ_{T_i}:

$$\mu_{T_i} = \frac{\sum_{i=1}^{N} T_i}{N}, \tag{8}$$

where *N* is the total number of agents.

The *mean reserve biomass* is the average of remaining reserve biomass on all occupied patches at the end of the simulation. The "occupied" condition is meant to exclude from the average the empty patches. The latter are not interesting as outcomes, because they all have the same vegetation dynamics and are in the same conditions, since they are not exploited (households only exploit patches that they occupy).

The *Gini-index* is a standard measure of inequality that takes values between 0 and 1. In the case of our model, it measures how the total livestock at the end of the simulation is distributed across the entire population of agents. A higher Gini-index represents larger differences in livestock numbers between the most endowed farmers and the least endowed.

The independent variable has **three factorial levels** and it represents whether the world is initialised with a homogeneous population of only E-RO agents, only E-LBD agents or only E-SL1 agents.

To control for the effects of other parameters of the model, we ran partial sensitivity analyses on the control variables indicated in Table 2, which in turn informed the final parameter values to be used as factorial levels in a full sensitivity analysis for assessing the robustness of the model results. In an iterative way, we started with very simple sub-DOEs for the behaviour of one agent and gradually added more complexity to explore the parameter space in a systematic manner and to ensure that our selection of final parameter values was robust.

Population size refers to the total number of agents with which the model world is initialised before the run. *Initial herd size* is the total number of livestock that each agent is endowed with at the beginning of the simulation. All agents start with the same initial herd size. *Model seed* is a control for the stochastic processes in the model to make the different treatments comparable when we have only few repetitions for each treatment. In our extension model, there are two sources of stochasticity: the distribution of *r-parameters* in the population and the spatial placement of the agents on the 2D map. As such, depending on the effect that we want to observe we need either a fixed seed for the various treatments, or a very large number of *repetitions* of the model run. The number of repetitions refers to how many times the exact same experimental treatment is run. For obvious reasons, repetitions only make sense when the model seed is not fixed. For the *r-parameter*, because it is drawn from a discrete uniform distribution centred around 0, the combination of a large number of agents with a large number of model runs (repetitions) means that the average *r-parameter* value in the model will converge to 0, thus making it possibly to compare model outputs on the dependent variables under conditions of stochasticity. Lastly, *time steps* defines the number of decision-rounds after which the model stops and outcomes are compared.

Based on our various sensitivity analyses, we settled in this experiment for the input parameters in Table 3. Parameters that were not specific to our extension module were set using the default values from RAGE (see Appendix).

Table 3. Input parameters for experiments with three different homogeneous populations of agent types (total runs: 3000)

Parameter name	Value	Description
Behavioural-type	*3 factorial levels:* E-RO, E-LBD, E-SL1	Agent type / learning process
Livestock-init	90	Initial herd size
HH-init (number-households)	50	Population size / initial number of households
Timesteps	100	Simulation length
	Random	Model Seed
	1000	Repetitions

Experiment 2

The second experiment was designed as a verification of the model and of the results from Experiment 1. We wanted to check how a heterogeneous population with equal proportions of three types of agents will perform on the three indicators of interest in comparison to the three homogeneous populations from the first experiment. Treatments and input parameters are detailed in Table 4. We expected that the mixed population will produce social-ecological outcomes that are in the middle of the outcomes produced by the three separate homogeneous populations.

Table 4. Treatments and input parameters for Experiment 2

Treatment	Initial herd size	Total number of households	Number of initial E-RO agents	Number of initial E-LBD agents	Number of initial E-SL1 agents	Model seed & Repetitions
Heterogeneous population	90	60	20	20	20	Fixed (1 run)
Homogeneous population E-RO only	90	60	60	-	-	Fixed (1 run)
Homogeneous population E-LBD only	90	60	-	60	-	Fixed (1 run)
Homogeneous population E-SL1 only	90	60	-	-	60	Fixed (1 run)

Experiment 3

In the final experiment we seek to find out how learning strategies interact within a mixed population and, in particular, we want to study the conditions under which learning-by-doing diffuses. Many research questions could be asked to characterise the effects of strategy switching, but we focus on learning-by-doing diffusion because, as we will show in the results, it has proven to be the more successful behavioural strategy in economic terms, and economic goals are the driver of strategy switching. As such, we ask: will learning-by-doing diffuse in a population even when it is difficult for others to find out about it?

We design an experiment where, for each agent type, we track the evolution of agent counts over time. The relevant parameters and treatments are detailed in Table 5. We keep the same value for the initial herd size (90) and we control for the total population size because household density on the map matters to whether other agents might find out about the E-LBD behaviour or not. At low total household density, the E-LBD households might be isolated spatially, so that no other agent can copy their learning type. The knowledge radius k is 1 for all treatments, meaning that the agents can observe other agents within their Moore neighbourhood only. We choose three levels for the total number of households (15, 45, 85) which correspond to low, middle and high household density in the world (10x10 patches). We then pick three very low initial numbers of E-LBD households, 1, 3 and 5, respectively, while the rest of the households are assigned equally to the E-RO and E-SL1 behaviour. Note that the same initial numbers of E-LBD households correspond to different proportions in the populations, depending on the total number of households. We run the model 100 times for each treatment.

Table 5. Treatments and parameter inputs for Experiment 3 (900 runs)

Treatment	Initial herd size	Total number of households	Number of initial E-LBD agents	Number of initial E-RO agents	Number of initial E-SL1 agents	Repetitions
1.	90	15	1	12	12	100
2.	90	15	3	6	6	100
3.	90	15	5	5	5	100
4.	90	45	1	22	22	100
5.	90	45	3	21	21	100
6.	90	45	5	20	20	100
7.	90	85	1	42	42	100
8.	90	85	3	41	41	100
9.	90	85	5	40	40	100

5. Results

5.1. Learning-by-doing vs. social learning farmers' decisions outcomes

We evaluate model outputs for three types of outcomes – economic, ecological and social – as explained previously.

Economically, our findings suggest that the learning-by-doing agents are more successful than other agent types at maintaining high numbers of livestock on their pastures without overshooting too often or too much the carrying capacity (Figure 7). The immediate adjustment of livestock numbers in response to observed decline in pasture state means that action is taken while the reserve biomass is still high, permitting a quick regeneration of the pasture condition and minimum economic losses (total number of livestock that are hungry).

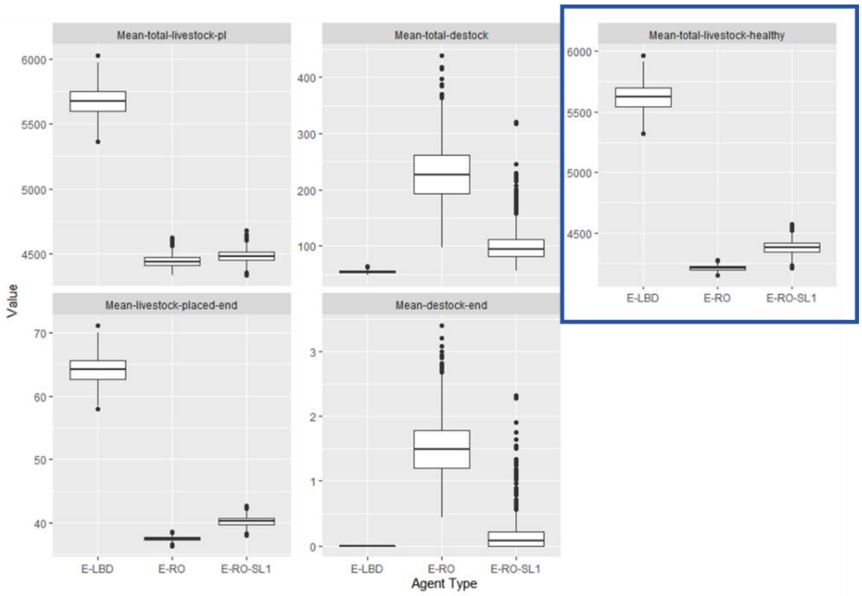

Figure 7. Economic outcomes for three homogeneous populations (Experiment 1)

Social learning behaviour also leads to higher numbers of healthy livestock than the baseline no-learning behaviour, but still much lower than learning-by-doing. Compared to no-learning agents, social learning agents make herd size decisions that are closer to the carrying capacity of the pasture, as indicated by a much lower number of hungry livestock (destock). That social learning has a relatively low economic success as a strategy appears to be consistent with Le et al. (2012) who also find that imitation does not improve incomes very much. The reason could be linked to the different conditions in which the successful farmers who are imitated are operating vs. those of the imitator. In the case of our model, by the time that a social learning agent starts imitating the attributes of a neighbour with more livestock, their pasture is already in a different condition, so they can no longer make up for lost economic opportunities if the final results are measured with a cumulative variable such as the ***total*** healthy livestock, i.e. over the entire simulation. However, this explanation is not very solid, as illustrated by the alternative indicators shown on the second row of Figure 7, where livestock numbers are measured at the end of the simulation. It is clear that the social learning agents manage to stabilize their herd size decision at a higher value than no learning agents (high livestock-placed-end), which is also better matching the pasture condition (low destock-end). An alternative explanation could be that certain ranges of the *r-parameter* lead to higher livestock numbers and, so, they end up being selected more often by social learners. This is a plausible explanation that warrants further study, but it is interesting to note that social learners' economic outcomes still remain much behind those of the learning-by-doing agents.

With respect to the **ecological** outcomes, social learning agents maintain a slightly higher level of reserve biomass than learning-by-doing agents **Error! Reference source not found.**. This is to be expected, given the economic outcomes discussed above, as smaller herd sizes translate into a lower grazing pressure on the pasture. A more interesting result is that the pastures occupied by no-learning agents have the lowest reserve biomass from all agent types at the end of the simulation. Although they have not placed very large herd sizes on their pastures, they degraded the pastures the most and they also had most livestock hungry. This shows that it is not the absolute numbers that are most important in the decisions made, but what matters is the extent to which the herd size decisions approximate the carrying capacity of the pasture. As the theory would suggest, both learning processes appear to contribute to reducing structural uncertainty about the "optimal" level of resource use by allowing learning agents to better match their decisions to the pasture state than baseline agents.

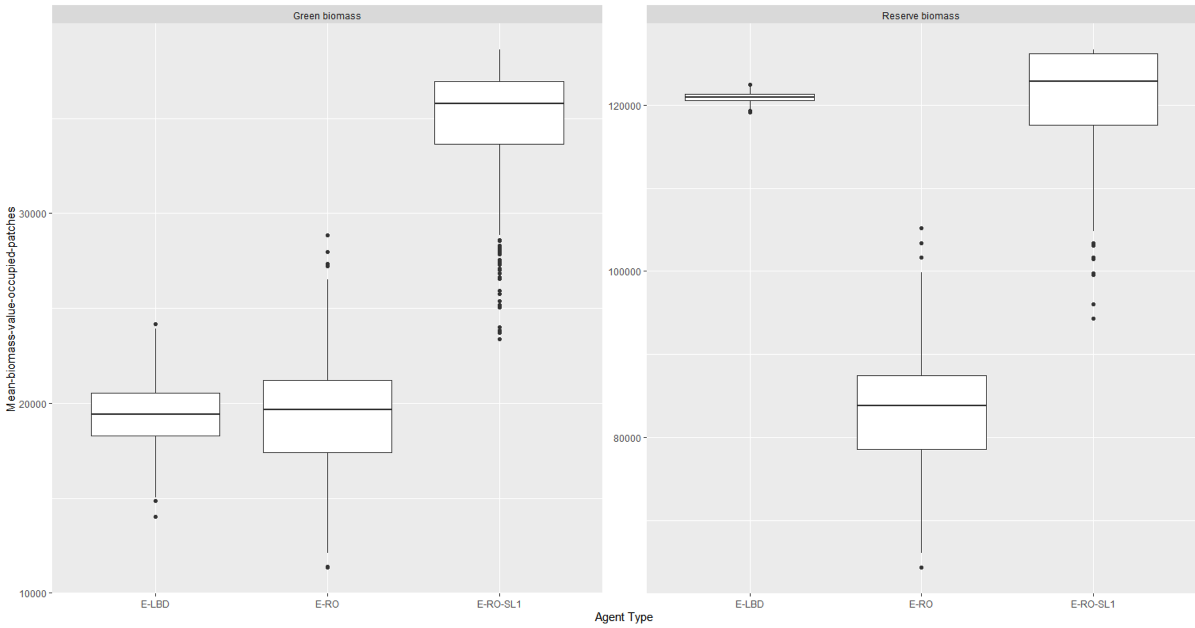

Figure 8. Ecological outcomes for three homogeneous populations (Experiment 1)

Finally, from a social outcomes perspective, learning-by-doing behaviour results in the highest economic inequality at the end of the simulation, with some agents having stabilized their herd size at low values and others at much higher values. A world of social learning agents also results in higher economic inequality than a world of non-learners, but the Gini-index is much lower than in the case of the learning-by-doing population. An explanation for this result is not straightforward, and of course no generalizations can be made from our very simple model, but it is interesting to speculate about the cultural role that imitation might have in evening out differences.

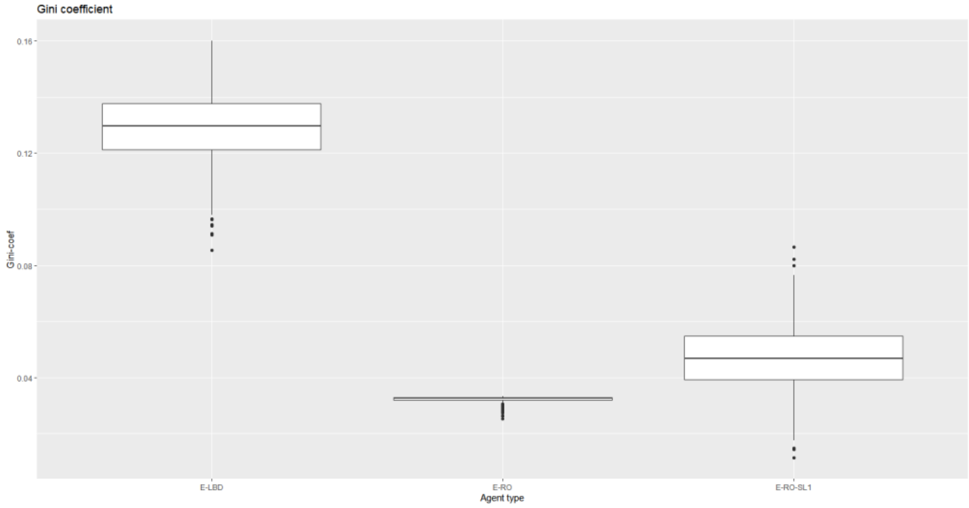

Figure 9. Social outcomes for three homogeneous populations (Experiment 1)

5.2. Learning interactions: Heterogeneous vs. homogeneous populations

The results of our second experiment correspond to our expectations and their main role is that they verify the correct functioning of our model. Verification refers to ensuring that the model is doing what the modeller expects it to do (Jakeman et al., 2006). However, we also wanted to check for any surprising emergent patterns that we may have not anticipated.

We here compare the economic, ecological and social outcomes of a mixed population of 20 E-RO, 20 E-LBD and 20 E-SL1 agents with three homogeneous populations of each type (Figure 10 a-c). E-HET is the mixed population treatment. As we hypothesised, the mean total livestock healthy of the mixed population is in the middle of the range delimited by the homogeneous treatments. The same is true for the mean reserve biomass values.

For the social outcomes, a heterogeneous population of learners leads to the highest economic inequality of all treatments, i.e. the highest differences in livestock herds among agents at the end of the simulation. This result is consistent with expectations, because the heterogeneous treatment introduces the herd values of the E-RO agents which are much lower than those of the "poorest" E-LBD farmers, which means that overall inequality in the population will be higher.

a)

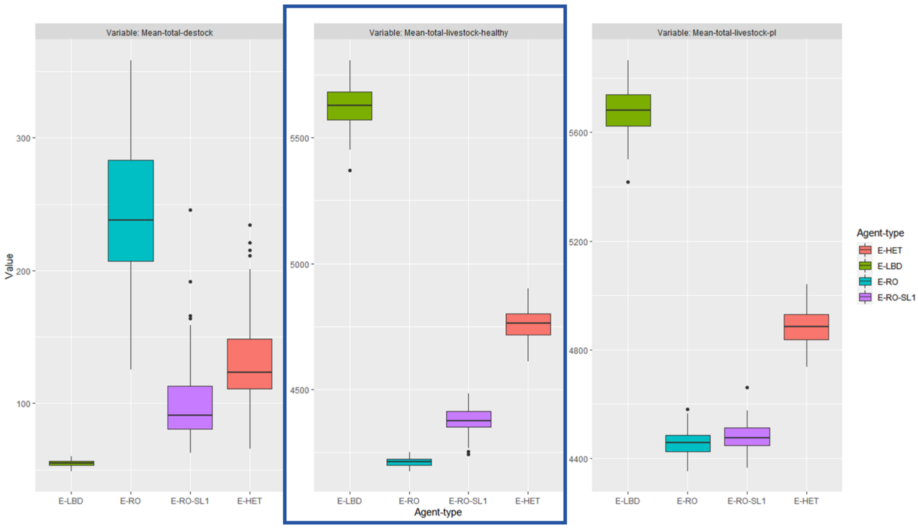

b)

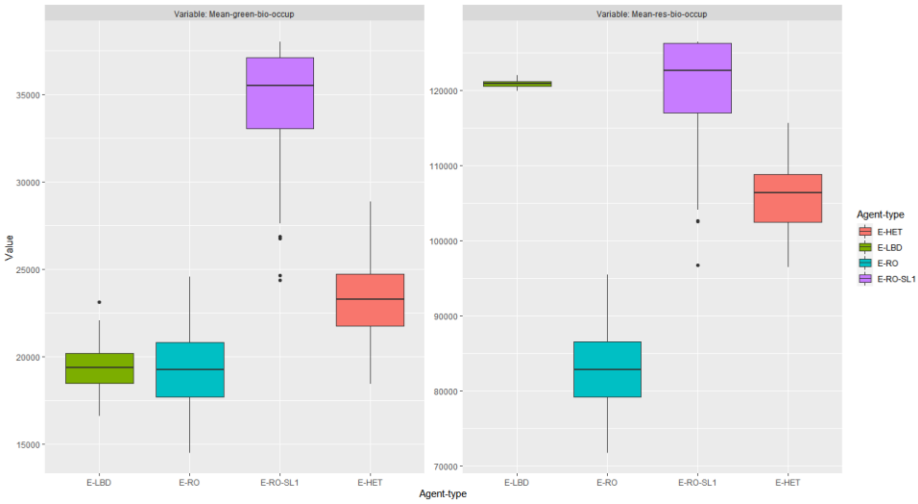

c)

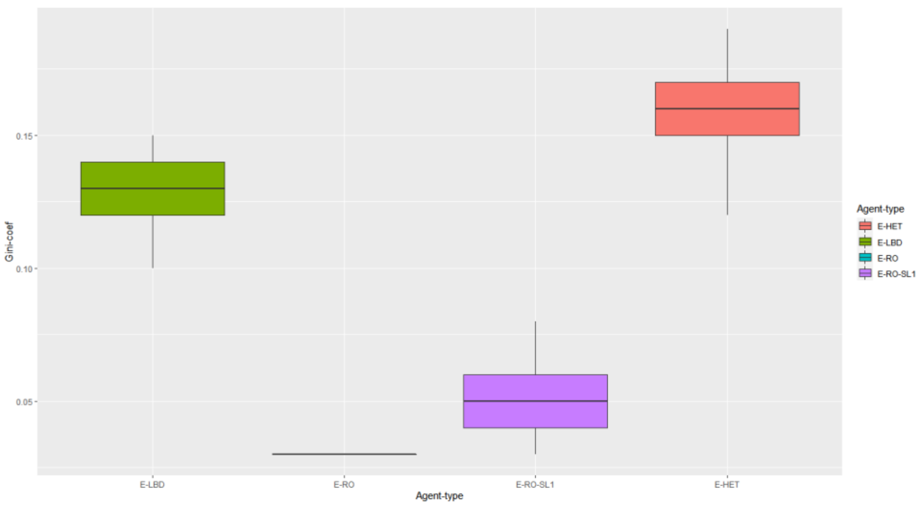

Figure 10. Outcomes for heterogeneous vs. homogeneous populations of learning agents. a) Economic outcomes; b) Ecological outcomes; c) Social outcomes.

An unanticipated result is depicted in Figure 11 when we look not at the average performance of all agents in the heterogeneous treatment, but rather when we compare the performance of various agent types depending on whether this was realised in a homogeneous vs. heterogeneous environment. We notice that, while Experiment 1 has shown that learning-by-doing agents hold pastures with a higher reserve biomass at the end of the simulation, this only holds true as long as the agents perform in a homogeneous environment; in the heterogeneous setting, the relationship is reversed. Social learners will have their pastures in a better condition when they will have been surrounded only by other social learners rather than by diverse learning types. The explanation for this is that the *r-parameter* selection bias discussed before might be even stronger in a population with higher economic inequality (driven by the presence of learning-by-doing agents) because the chances increase to observe a neighbour with a bigger herd size than one's own.

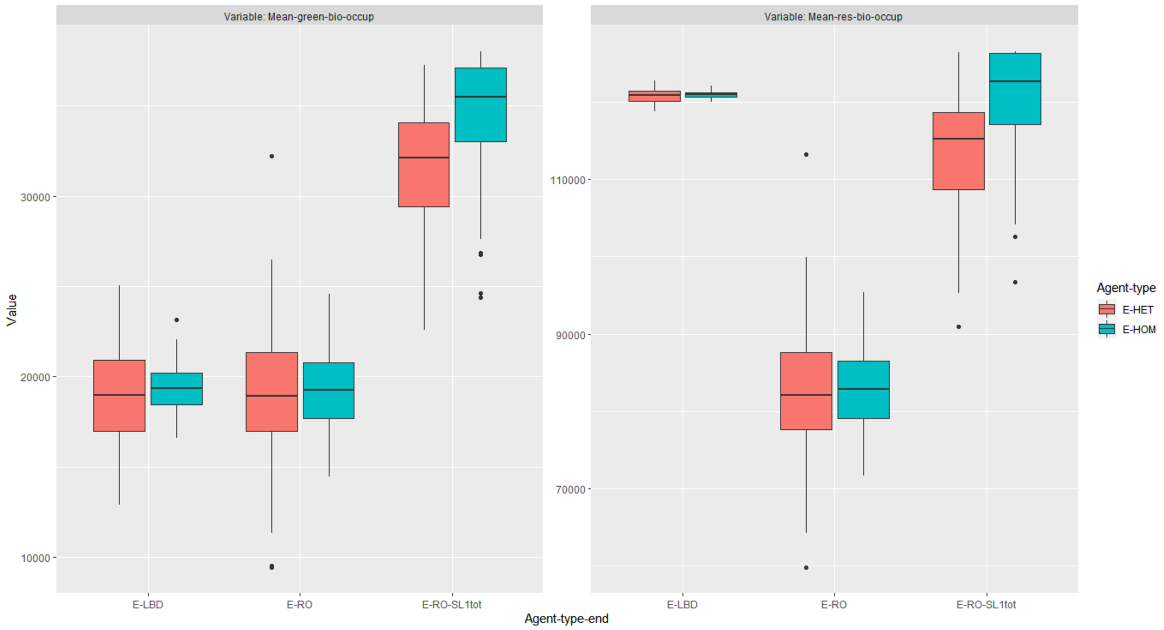

Figure 11. Ecological outcomes for specific agent types when embedded in a heterogeneous vs. a homogeneous population.

5.3. Learning interactions: Strategy switching

The results of our third experiment on the diffusion of the learning-by-doing behaviour suggest that a strategy that is perceived as successful (E-LBD) will spread even if the chances of it to be encountered are very small. The first column of Figure 12 shows that even at high population densities, one agent will not be able to spread their behaviour very efficiently, and they will only reach very small corners of the network (on average 5 agents in a population of 85, from 1.2% to 6%). More interestingly, in a low and medium-density environment, 3 agents succeed in spreading their behaviour to between 22% and 27% of all agents reaching an almost equal share of agents with the other behaviours. With 5 agents, the behaviour spreads very quickly and it dominates the world in a relatively low- or medium-density environment. In other words, there are tipping points in the diffusion dynamics and successful behaviours might require much less of an initial crowd to become dominant than one might assume. For instance, panel 6 of Figure 12 shows that having an initial share of 11% E-LBD agents is already enough for the behaviour to become dominant. While these numbers cannot be generalized beyond the limited context of this model, our results are consistent with studies of critical mass that have demonstrated that tipping points in social conventions might occur when as little as 10% of the population engages committedly in a specific behaviour (see e.g. Centola et al., 2018).

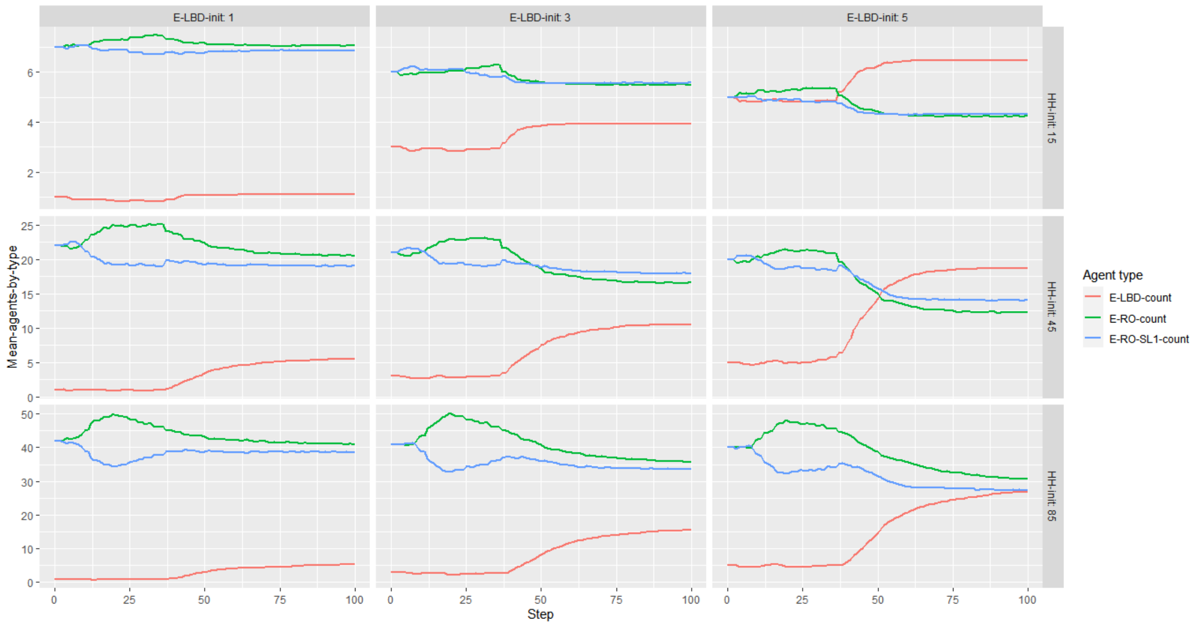

Figure 12. Evolution of agent counts over time, by agent type – means for 100 repetitions. Results are displayed for 9 different treatments. Each column corresponds to a different initial number of E-LBD agents, while each row corresponds to a different initial total number of households.

6. Discussion and Conclusion

Our work sought to contribute to the broader agenda of modelling and understanding adaptive decision-making by studying the role of learning processes in the specific context of smallholder farmers' behaviour. In particular, we set to: a) make a conceptual and methodological contribution as to what learning-by-doing and social learning are and how these processes might be represented in social-ecological models; and b) advance our theoretical understanding of how learning-by-doing and social learning interact and affect smallholder farmers' decision outcomes.

Although there is a long tradition of modelling learning, especially within economics, the main concept used has been that of a learning curve, which is very limited in its ability to represent realistic, non-rational-choice human behaviour, as well as to capture the diversity of settings and levels at which learning happens. Learning processes, as discussed, for instance, within adaptive (co)-management, require other conceptualisations if they are to be used in models. Within sustainability science and social-ecological systems research, as well as in most agent-based agricultural models, learning processes are seldom represented. One possible reason for this is that there is little agreement and clarity of what terms such as "learning-by-doing" or "social learning" mean in these contexts, and there is very little theoretical ground to build upon. That is why, our first objective was to work towards clarifying the concepts, by explicating who learns what and to what effects. Preliminary ideas towards a conceptual framework for modelling learning-by-doing and social learning are presented in Section 3 and, in themselves, they amount to an original contribution of our work.

For our second objective, we developed a module with learning agents which extends a pre-existing social-ecological rangeland grazing ABM, RAGE (Dressler et al., 2019). We implemented three types of agents, baseline, non-learning agents (E-RO), learning-by-doing agents (E-LBD) and social learning agents (E-SL1). We also implemented a meta-type of social learning as strategy switching, which allows all agents to change their agent type / learning process under certain conditions. We then ran three different experiments, following a DOE approach, to answer our research questions.

First, we analysed and compared the social-ecological performance of homogeneous populations of each of the three types of agents. Our results suggest that learning-by-doing behaviour maximises economic performance, while social learning (1) maximises environmental performance (understood here as minimum pressure on the pasture). Most importantly, in line with what theory would suggest, both learning processes appear to contribute to reducing structural uncertainty about the "optimal" level of resource use, by allowing learning agents to better match their decisions to the pasture state than baseline agents.

Second, we compared outcomes in a heterogeneous world of learning agents with outcomes in homogeneous worlds. As expected, heterogeneous environments with equal shares of three types of agents led to global outcomes that are the average of the three homogeneous environments corresponding to each agent type. A surprising result, however, was that social learning behaviour resulted in better ecological outcomes when agents where deciding in a homogeneous rather than heterogeneous environment. This finding warrants further exploration.

Third, we analysed strategy switching with a focus on the adoption of learning-by-doing. Agents decide whether to switch strategies based on economic considerations. Because learning-by-doing performs better economically (as shown in our first experiment), it is perceived by agents as the more successful strategy, hence our choice to focus on E-LBD's diffusion dynamics. We were particularly interested to see how learning-by-doing behaviour spreads in worlds initialised with only few agents of this type, i.e. where the initial chances of others to encounter the learning-by-doing behaviour are very small. Our findings show that as few as 3 agents are enough to spread, over 100 time steps, the behaviour to at least 20% of the population, regardless of the total size of the population. Moreover, an initial share of 11% E-LBD agents in the total population is enough for crossing tipping points so that the learning-by-doing behaviour will become dominant in the population at the end of the simulation. This result is consistent with other studies of critical mass that have also found tipping points in network contagions at values ranging between 10 and 40% (Centola et al., 2018).

The main limitations of our model are those related to it being a stylised ABM, which means that it is not trying to reproduce an empirical situation in great detail. In addition, the theoretical gap in conceptualising and operationalising learning-by-doing and social learning for modelling means that we still had to make many micro-choices in our implementation, which could be debated. Our approach in dealing with this limitation has been to clearly define an operationalisation framework for our concepts, to explain the methodological choices we made and to suggest alternatives for implementation which could be used by future models trying to replicate our overall results. For instance, although our model tries to step away from classical rational choice models and to introduce behavioural heterogeneity and specific parameters for dealing with uncertainties (r-parameter), certain aspects of agent behaviour still relied on economic utility maximisation. Janssen (2016) wrote about implementing rational choice agents: "In doing this exercise one realizes that a pure rational actor approach is difficult to achieve in agent-based models" (p.1693). Similarly, we believe that implementing agents that completely step away from rational choice is difficult and perhaps unrealistic. What we think a useful further theoretical development would be is for learning processes to be clearly mapped and operationalised on those meta theories of human behaviour that match in terms of ontological and epistemological assumptions (see our discussion in Section 2.4).

With our study we are only starting to scratch the surface of understanding the role of learning in farmers' decision-making and much more conceptual and modelling work is needed to develop a "science of adaptive agents", in the words of Miller et al. (2007). Future research could consider implementing alternative, but comparable, operationalisations of learning-by-doing and social learning to see, similarly to a sensitivity analysis, if our results can be reproduced. In addition, developing ABM modules with a broader repertoire of learning processes, each with clearly specified assumptions and methodological choices, could be very useful for empirically-based models. A point to note is that, in our evaluation of learning outcomes, we use multiple indicators corresponding to three different areas of impact (economic, ecological and social) and that the learning processes in our model show mixed results across all three areas. Future research might try to combine these indicators into standardized units, so as to allow for multi-criteria optimisation if specific outcomes are desired. This would then open the avenue for such learning models to be used as policy tools in studying which type of learning should be encouraged by a social planner trying to achieve specific results as quickly as possible. Last, but not least, our model could, with minimum additional work, also be used to study institutional emergence / norm formation and collective action within social-ecological systems, an area which is also currently under-explored in AMBs, but very relevant from a policy perspective.

References

Apetrei, C.I., Caniglia, G., Wehrden, H. von, Lang, D.J., 2021. Just another buzzword? A systematic literature review of knowledge-related concepts in sustainability science. Global Environmental Change 68, 102222.

Argyris, C., Schön, D.A., 1978. Organizational learning (in eng). Addison-Wesley Pub. Co, Reading, Mass, 344 pp.

Arifovic, J., Ledyard, J., 2004. Scaling Up Learning Models in Public Good Games. J Public Economic Theory 6 (2), 203–238.

Armitage, D., Marschke, M., Plummer, R., 2008. Adaptive co-management and the paradox of learning (in en). Global Environmental Change 18 (1), 86–98.

Baird, J., Plummer, R., Haug, C., Huitema, D., 2014. Learning effects of interactive decision-making processes for climate change adaptation. Global Environmental Change 27, 51–63.

Bandura, A., 1977. Social learning theory. Prentice-Hall, Englewood Cliffs, London.

Beers, P.J., van Mierlo, B., Hoes, A.-C., 2016. Toward an Integrative Perspective on Social Learning in System Innovation Initiatives. E&S 21 (1).

Bennett, C.J., Howlett, M., 1992. The lessons of learning: Reconciling theories of policy learning and policy change. Policy Sci 25 (3), 275–294.

Berglund, C., Söderholm, P., 2006. Modeling technical change in energy system analysis: analyzing the introduction of learning-by-doing in bottom-up energy models. Energy Policy 34 (12), 1344–1356.

Biggs, R., Clements, H., Vos, A. de, Folke, C., Manyani, A., Maciejewski, K., Mart\'\in-López, B., Preiser, R., Selomane, O., Schlüter, M., 2021. What are social-ecological systems and social-ecological systems research? In: Biggs, R., Vos, A. de, Preiser, R., Clements, H., Maciejewski, K., Schlüter, M. (Eds.), The Routledge Handbook of Research Methods for Social-Ecological Systems. Taylor & Francis, pp. 3–26.

Bointner, R., Schubert, K., 2016. The influence of experience and knowledge on reactor safety in Germany. Journal of Nuclear Science and Technology 53 (7), 1009–1020.

Bousquet, F., Le Page, C., 2004. Multi-agent simulations and ecosystem management: a review. Ecological Modelling 176 (3-4), 313–332.

Brenner, T., 2006. Chapter 18 Agent Learning Representation: Advice on Modelling Economic Learning. In: Tesfatsion, L., Judd, K.L. (Eds.), Handbook of Computational Economics. Agent-Based Computational Economics, vol. 2. Elsevier, Burlington, pp. 895–947.

Brown, C., Alexander, P., Holzhauer, S., Rounsevell, M.D.A., 2017. Behavioral models of climate change adaptation and mitigation in land-based sectors. WIREs Clim Change 8 (2), e448.

Camerer, C.F., 1995. Individual Decision-Making. In: Kagel, J.H., Roth, A.E. (Eds.), The Handbook of Experimental Economics. Princeton University Press, pp. 587–673.

Centola, D., Becker, J., Brackbill, D., Baronchelli, A., 2018. Experimental evidence for tipping points in social convention (in eng). Science (New York, N.Y.) 360 (6393), 1116–1119.

Darnhofer, I., Bellon, S., Dedieu, B., Milestad, R., 2010. Adaptiveness to enhance the sustainability of farming systems. A review. Agron. Sustain. Dev. 30 (3), 545–555.

Dewey, J., 1988. Experience and Nature. The Later Works of John Dewey, 1925-1953. Southern Illinois University Press, Carbondale & Edwardsville.

Diduck, A., 2010. The Learning Dimension of Adaptive Capacity: Untangling the Multi-level Connections. In: Armitage, D., Plummer, R. (Eds.), Adaptive Capacity and Environmental Governance. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 199–221.

Dosi, G., Grazzi, M., Mathew, N., 2017. The cost-quantity relations and the diverse patterns of "learning by doing": Evidence from India. Research Policy 46 (10), 1873–1886.

Dressler, G., Groeneveld, J., Buchmann, C.M., Guo, C., Hase, N., Thober, J., Frank, K., Müller, B., 2019. Implications of behavioral change for the resilience of pastoral systems—Lessons from an agent-based model. Ecological Complexity 40, 100710.
Filatova, T., Verburg, P.H., Parker, D.C., Stannard, C.A., 2013. Spatial agent-based models for socio-ecological systems: Challenges and prospects. Environmental Modelling & Software 45, 1–7.
Foley, J.A., Defries, R., Asner, G.P., Barford, C., Bonan, G., Carpenter, S.R., Chapin, F.S., Coe, M.T., Daily, G.C., Gibbs, H.K., Helkowski, J.H., Holloway, T., Howard, E.A., Kucharik, C.J., Monfreda, C., Patz, J.A., Prentice, I.C., Ramankutty, N., Snyder, P.K., 2005. Global consequences of land use (in eng). Science (New York, N.Y.) 309 (5734), 570–574.
Glock, C.H., Grosse, E.H., Jaber, M.Y., Smunt, T.L., 2019. Applications of learning curves in production and operations management: A systematic literature review. Computers & Industrial Engineering 131, 422–441.
Groeneveld, J., Müller, B., Buchmann, C.M., Dressler, G., Guo, C., Hase, N., Hoffmann, F., John, F., Klassert, C., Lauf, T., Liebelt, V., Nolzen, H., Pannicke, N., Schulze, J., Weise, H., Schwarz, N., 2017. Theoretical foundations of human decision-making in agent-based land use models – A review. Environmental Modelling & Software 87, 39–48.
Huber, R., Bakker, M., Balmann, A., Berger, T., Bithell, M., Brown, C., Grêt-Regamey, A., Xiong, H., Le, Q.B., Mack, G., Meyfroidt, P., Millington, J., Müller, B., Polhill, J.G., Sun, Z., Seidl, R., Troost, C., Finger, R., 2018. Representation of decision-making in European agricultural agent-based models. Agricultural Systems 167, 143–160.
Hunecke, C., Engler, A., Jara-Rojas, R., Poortvliet, P.M., 2017. Understanding the role of social capital in adoption decisions: An application to irrigation technology. Agricultural Systems 153, 221–231.
IFAD, 2003. Smallholders, food security, and the environment. International Fund for Agricultural Development. https://www.ifad.org/documents/38714170/39135645/smallholders_report.pdf/133e8903-0204-4e7d-a780-bca847933f2e (accessed September 2021).
Jakeman, A.J., Letcher, R.A., Norton, J.P., 2006. Ten iterative steps in development and evaluation of environmental models. Environmental Modelling & Software 21 (5), 602–614.
Janssen, M.A., 2016. Impact of diverse behavioral theories on environmental management: explorations with Daisyworld. In: 2016 Winter Simulation Conference (WSC), pp. 1690–1701.
Janssen, S., van Ittersum, M.K., 2007. Assessing farm innovations and responses to policies: A review of bio-economic farm models. Agricultural Systems 94 (3), 622–636.
Jouvet, P.-A., Schumacher, I., 2012. Learning-by-doing and the costs of a backstop for energy transition and sustainability. Ecological Economics 73, 122–132.
Kahneman, D., Tversky, A. (Eds.), 2000. Choices, values, and frames (in eng). Cambridge University Press, Cambridge, 848 pp.
Kolb, D.A., 1984. Experiential learning: Experience as the source of learning and development. Prentice-Hall, Englewood Cliffs, N.J., 256 pp.
Le, Q.B., Seidl, R., Scholz, R.W., 2012. Feedback loops and types of adaptation in the modelling of land-use decisions in an agent-based simulation. Environmental Modelling & Software 27-28, 83–96.
Lindkvist, E., Norberg, J., 2014. Modeling experiential learning: The challenges posed by threshold dynamics for sustainable renewable resource management. Ecological Economics 104, 107–118.
Lorscheid, I., Heine, B.-O., Meyer, M., 2012. Opening the 'black box' of simulations: increased transparency and effective communication through the systematic design of experiments. Comput Math Organ Theory 18 (1), 22–62.
Marra, M., Pannell, D.J., Ghadim, A.A., 2003. The economics of risk, uncertainty and learning in the adoption of new agricultural technologies: where are we on the learning curve? Agricultural Systems 75 (2-3), 215–234.

Meyfroidt, P., 2013. Environmental cognitions, land change, and social–ecological feedbacks: an overview. Journal of land use science 8 (3), 341–367.
Miettinen, R., 2000. The concept of experiential learning and John Dewey's theory of reflective thought and action. International Journal of Lifelong Education 19 (1), 54–72.
Miketa, A., Schrattenholzer, L., 2004. Experiments with a methodology to model the role of R&D expenditures in energy technology learning processes; first results. Energy Policy 32 (15), 1679–1692.
Miller, J.H., Page, S.E., 2007. Complex Adaptive Systems: An Introduction to Computational Models of Social Life. Princeton University Press.
Miller, N.E., Dollard, J., 1941. Social learning and imitation.
Muelder, H., Filatova, T., 2018. One Theory - Many Formalizations: Testing Different Code Implementations of the Theory of Planned Behaviour in Energy Agent-Based Models. JASSS 21 (4).
Munaretto, S., Huitema, D., 2012. Adaptive Comanagement in the Venice Lagoon? An Analysis of Current Water and Environmental Management Practices and Prospects for Change. E&S 17 (2).
Muro, M., Jeffrey, P., 2008. A critical review of the theory and application of social learning in participatory natural resource management processes. Journal of Environmental Planning and Management 51 (3), 325–344.
Nhim, T., 2018. Does the interaction between public goods and common pool resources improve cooperation between communities? Young Scientists Summer Program (YSSP) Report. International Institute for Applied Systems Analysis (IIASA), Laxenburg, Austria.
Nowak, S., Matthews, L., Parker, A., 2017. A General Agent-Based Model of Social Learning. RAND Corporation.
Pahl-Wostl, C., 2009. A conceptual framework for analysing adaptive capacity and multi-level learning processes in resource governance regimes (in en). Global Environmental Change 19 (3), 354–365.
Parker, D.C., Manson, S.M., Janssen, M.A., Hoffmann, M.J., Deadman, P., 2003. Multi-Agent Systems for the Simulation of Land-Use and Land-Cover Change: A Review. Annals of the Association of American Geographers 93 (2), 314–337.
Preiser, R., Biggs, R., Vos, A. de, Folke, C., 2018. Social-ecological systems as complex adaptive systems: organizing principles for advancing research methods and approaches. E&S 23 (4).
Rebaudo, F., Dangles, O., 2011. Coupled information diffusion--pest dynamics models predict delayed benefits of farmer cooperation in pest management programs (in eng). PLoS computational biology 7 (10), e1002222.
Reed, M., Evely, A., Cundill, G.N.R., Fazey, I., Glass, J., Laing, A., Newig, J., Parrish, B., Prell, C., Raymond, C., Stringer, L.C., 2010. What is social learning? E&S 15 (4).
Regan, H.M., Colyvan, M., Burgman, M.A., 2002. A Taxonomy and Treatment of Uncertainty for Ecology and Conservation Biology. Ecological Applications 12 (2), 618–628.
Robert, M., Thomas, A., Bergez, J.-E., 2016. Processes of adaptation in farm decision-making models. A review. Agron. Sustain. Dev. 36 (4).
Rodela, R., 2011. Social Learning and Natural Resource Management: The Emergence of Three Research Perspectives. E&S 16 (4).
Rounsevell, M.D.A., Arneth, A., Alexander, P., Brown, D.G., Noblet-Ducoudré, N. de, Ellis, E., Finnigan, J., Galvin, K., Grigg, N., Harman, I., Lennox, J., Magliocca, N., Parker, D., O'Neill, B.C., Verburg, P.H., Young, O., 2014. Towards decision-based global land use models for improved understanding of the Earth system. Earth Syst. Dynam. 5 (1), 117–137.
Schlüter, M., Baeza, A., Dressler, G., Frank, K., Groeneveld, J., Jager, W., Janssen, M.A., McAllister, R.R.J., Müller, B., Orach, K., Schwarz, N., Wijermans, N., 2017. A framework for mapping and comparing behavioural theories in models of social-ecological systems. Ecological Economics 131, 21–35.

Schlüter, M., Haider, L.J., Lade, S.J., Lindkvist, E., Martin, R., Orach, K., Wijermans, N., Folke, C., 2019. Capturing emergent phenomena in social-ecological systems: an analytical framework. E&S 24 (3).

Schlüter, M., McAllister, R.R.J., Arlinghaus, R., BUNNEFELD, N., Eisenack, K., Hölker, F., Milner-Gulland, E.J., Müller, B., Nicholson, E., Quaas, M., Stöven, M., 2012. New Horizon for Managing the Environment: A Review of Coupled Social-Ecological Systems Modeling. Natural Resource Modeling 25 (1), 219–272.

Scholz, R.W., 2011. Environmental Literacy in Science and Society: From Knowledge to Decisions (in eng). Cambridge Univ. Press, Cambridge, 631 pp.

Schulze, J., Müller, B., Groeneveld, J., Grimm, V., 2017. Agent-Based Modelling of Social-Ecological Systems: Achievements, Challenges, and a Way Forward. JASSS 20 (2).

Schusler, T.M., Decker, D.J., Pfeffer, M.J.A.X., 2003. Social Learning for Collaborative Natural Resource Management. Society & Natural Resources 16 (4), 309–326.

Schwarz, N., Dressler, G., Frank, K., Jager, W., Janssen, M., Müller, B., Schlüter, M., Wijermans, N., Groeneveld, J., 2019. Formalising theories of human decision-making for agent-based modelling of social-ecological systems: practical lessons learned and ways forward. SESMO 2, 16340.

Simon, H.A., 1955. A Behavioral Model of Rational Choice (in af). The Quarterly Journal of Economics 69 (1), 99.

Skinner, B.F., 1938. The Behavior of Organisms (New York: Appleton).

Thompson, P., 2011. Learning by Doing. In: Hall, B.H., Rosenberg, N. (Eds.), Handbook of the economics of innovation. Volume 1, vol. 1, 1. ed., reprinted. ed. North Holland, Amsterdam, pp. 429–476.

Tilman, D., Fargione, J., Wolff, B., D'Antonio, C., Dobson, A., Howarth, R., Schindler, D., Schlesinger, W.H., Simberloff, D., Swackhamer, D., 2001. Forecasting agriculturally driven global environmental change (in eng). Science (New York, N.Y.) 292 (5515), 281–284.

Williams, B.K., Brown, E.D., 2016. Technical challenges in the application of adaptive management. Biological Conservation 195, 255–263.

Wright, T.P., 1936. Factors Affecting the Cost of Airplanes. Journal of the Aeronautical Sciences 3 (4), 122–128.

Appendix: Supplementary Information

Parameter values

Table S1. Default parameter values from RAGE that were used in all simulations

Parameter name	Value	Description
intake	640	Amount of fodder needed per livestock unit
gr1	0.5	Support parameter vegetation submodel
gr2	0.1	Support parameter vegetation submodel
w	0.8	Support parameter vegetation submodel
mg	0.1	Support parameter vegetation submodel
mr	0.05	Support parameter vegetation submodel
lambda	0.5	Support parameter vegetation submodel
Rmax	150000	Support parameter vegetation submodel
d	1/Rmax	Support parameter vegetation submodel
R0part	0.6	Support parameter vegetation submodel

Error variance analysis

Low values: HH-init 50, Livestock-init 90

E-RO:

	Repetitions:	100	500	1000	5000
Gini-coef					
	MEAN	0.03	0.03	0.03	0.03
	SD	0	0	0	0
	Coeff-var	0	0	0	0
Mean-green-bio					
	MEAN	53595.19	53661.65	53632.75	53599.4
	SD	1561.66	1562.72	1606.43	1576.29
	Coeff-var	0.029	0.029	0.03	0.029
Mean-res-bio					
	MEAN	107188.79	107321.72	107263.91	107197.2
	SD	3123.55	3125.68	3213.09	3152.81
	Coeff-var	0.029	0.029	0.03	0.029
Mean-total-destock					
	MEAN	227.7	235.43	230.25	232.59
	SD	60.74	55.88	51.28	55.08
	Coeff-var	0.267	0.237	0.223	0.237
Mean-total-livestock-healthy					
	MEAN	4211.56	4212.63	4212.34	4211.91
	SD	20.15	19.85	20.33	19.97
	Coeff-var	0.005	0.005	0.005	0.005
Mean-total-livestock-pl					
	MEAN	4439.26	4448.06	4442.59	4444.5
	SD	52.43	48.17	45.47	48.45
	Coeff-var	0.012	0.011	0.01	0.011

E-LBD:

	Repetitions:	100	500	1000	5000
Gini-coef					
	MEAN	0.13	0.13	0.13	0.13
	SD	0.01	0.01	0.01	0.01
	Coeff-var	0.077	0.077	0.077	0.077
Mean-green-bio					
	MEAN	63062.69	63038.79	63045.02	63049.52
	SD	137.6	134.83	138.86	139.06
	Coeff-var	0.002	0.002	0.002	0.002
Mean-res-bio					
	MEAN	126127.4	126079.7	126092.1	126101.1
	SD	274.91	269.37	277.42	277.82
	Coeff-var	0.002	0.002	0.002	0.002
Mean-total-destock					
	MEAN	54.22	54.62	54.54	54.44
	SD	2.89	2.77	2.84	2.85
	Coeff-var	0.053	0.051	0.052	0.052
Mean-total-livestock-healthy					
	MEAN	5606.06	5677.37	5673.67	5615.64
	SD	109.21	108.71	110.98	108.54
	Coeff-var	0.019	0.019	0.02	0.019
Mean-total-livestock-pl					
	MEAN	5660.29	5622.75	5619.14	5670.08
	SD	112.01	106.04	108.26	111.28
	Coeff-var	0.02	0.019	0.019	0.02

Sensitivity analyses

1. **One household – sensitivity analysis for the effect of initial livestock numbers and different values of the r-parameter on economic outcomes**

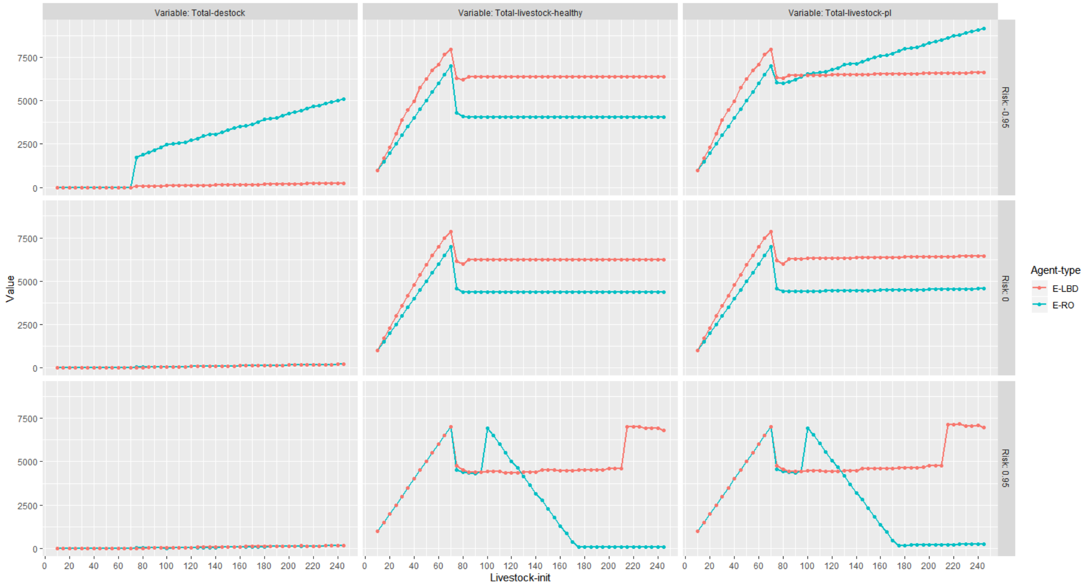

2. **One household – sensitivity analysis for the effect of initial livestock numbers and different values of the r-parameter on reserve biomass and green biomass**

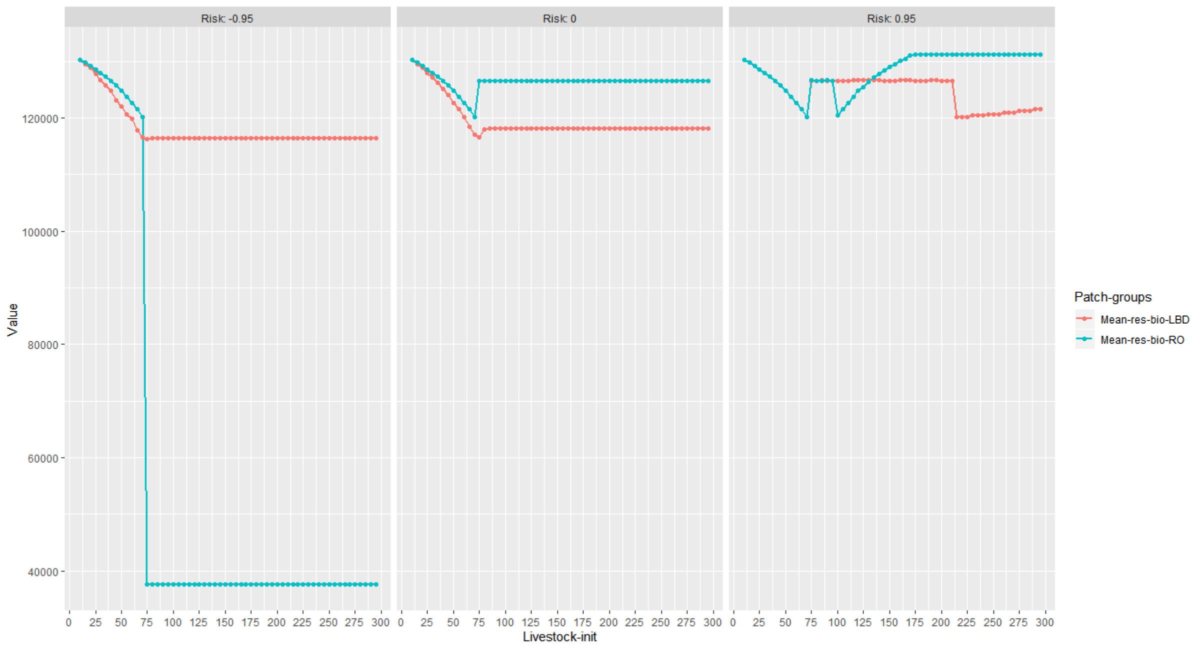

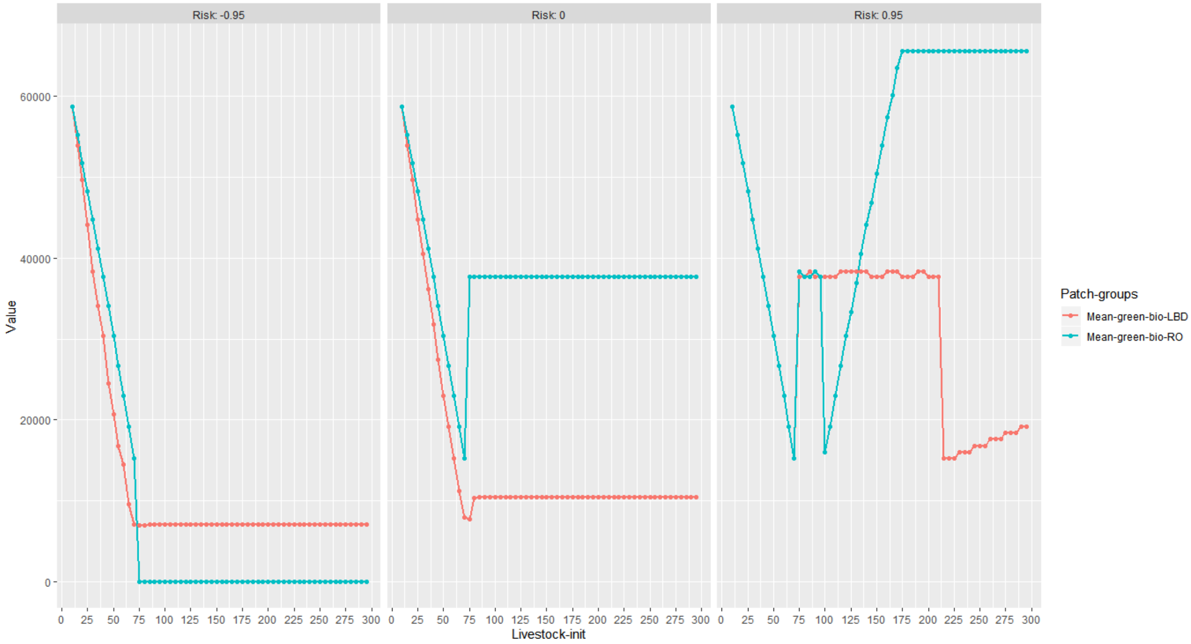

3. **Multiple households, livestock-init=80 – sensitivity analysis of the effect of the initial number of households (HH-init) on economic and ecological outcomes by agent type**

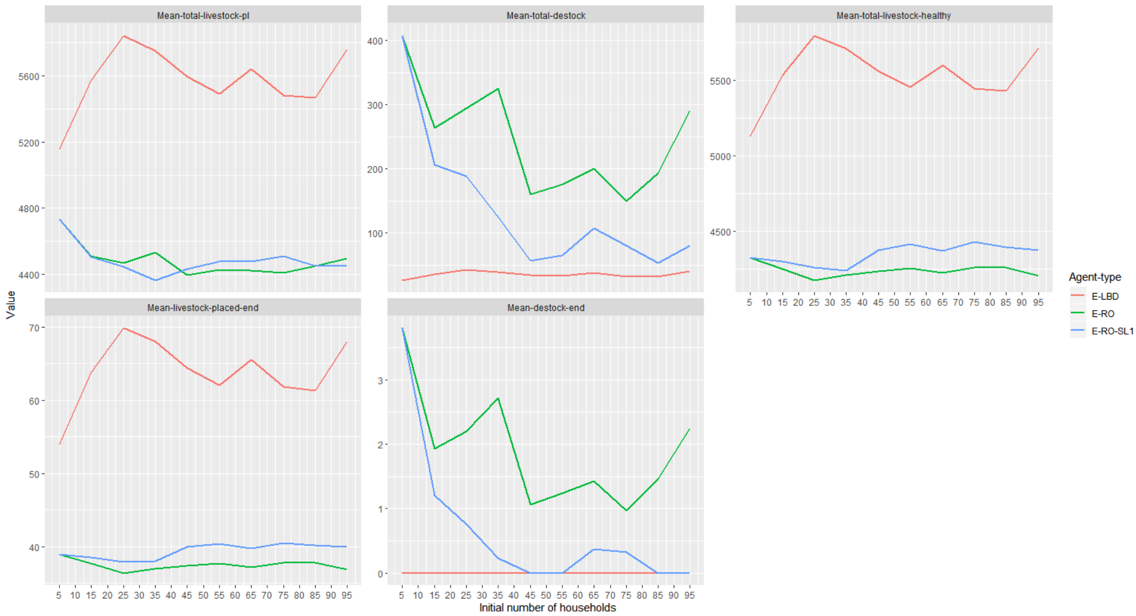

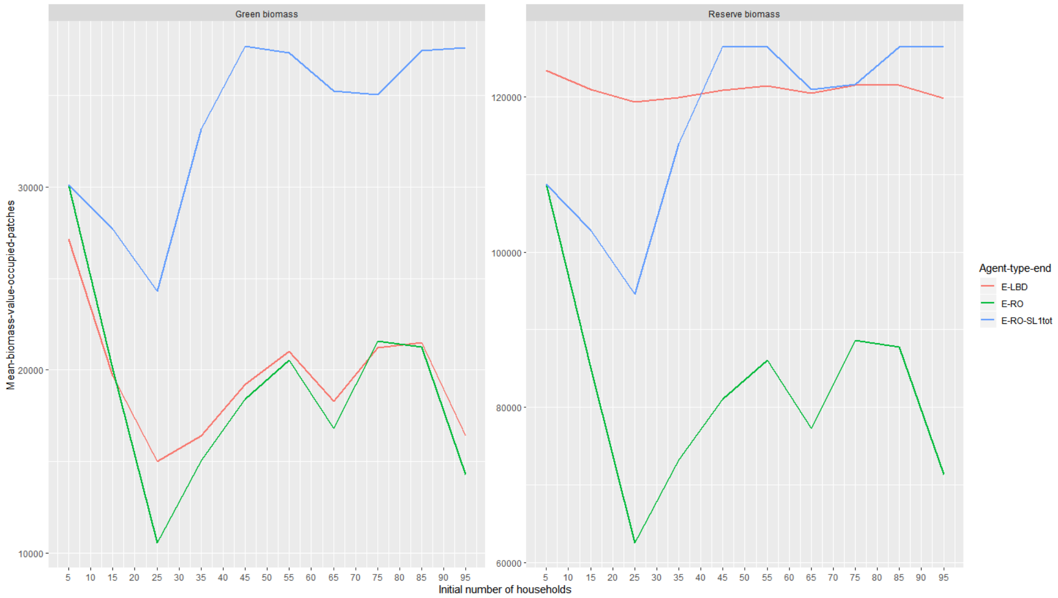

4. Sensitivity analysis for initial mean R parameter in the population and its effects on biomass

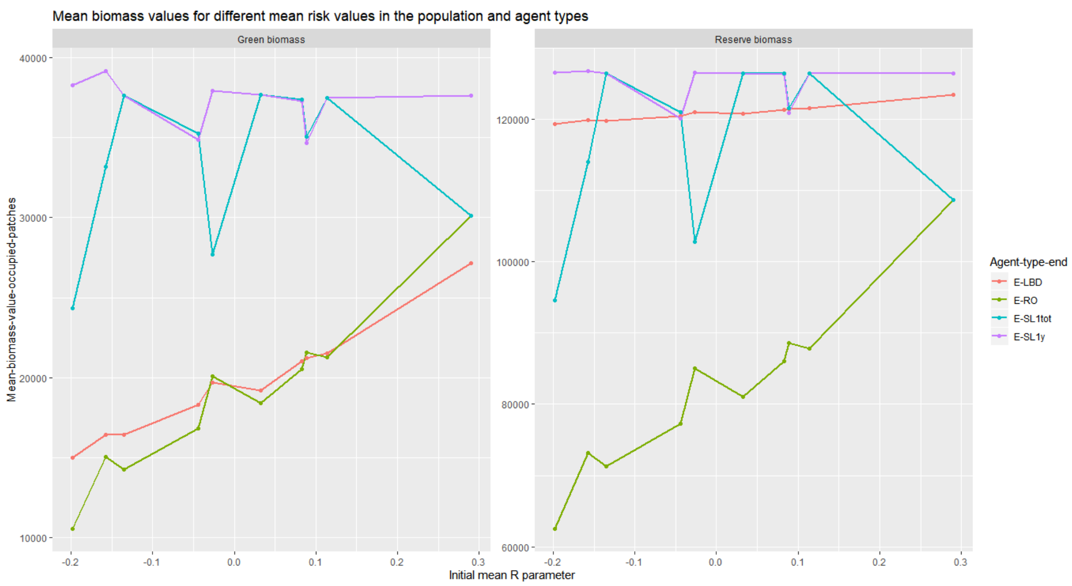

5. Full sensitivity analysis with 100 repetitions (9000 runs):

Factor levels were chosen based on sensitivity analyses 1-4 above:

["behavioral-type" "E-LBD" "E-RO" "E-RO-SL1"]
["number-households" 5 30 50 70 95]
["livestock-init" 70 75 90 110 160 250]
["timesteps" 100]
["homog-behav-types?" true]
["knowledge-radius" 1]
["extension-model?" true]

Total livestock healthy:

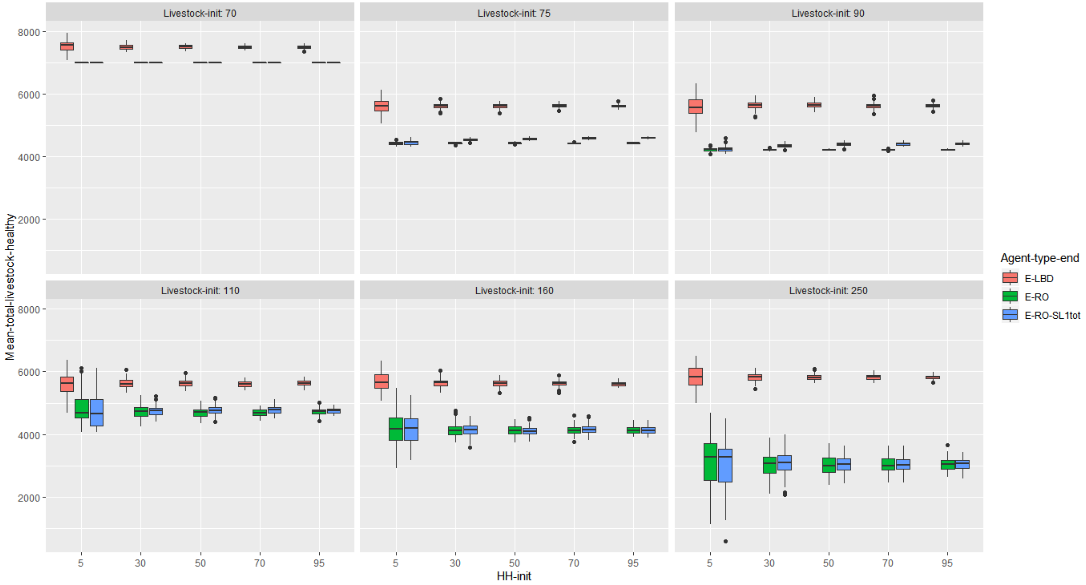

Total livestock healthy with splitting of social learning results between E-SL1 agents who did learn (SL1y) by the end of the simulation and those who did not (SL1n):

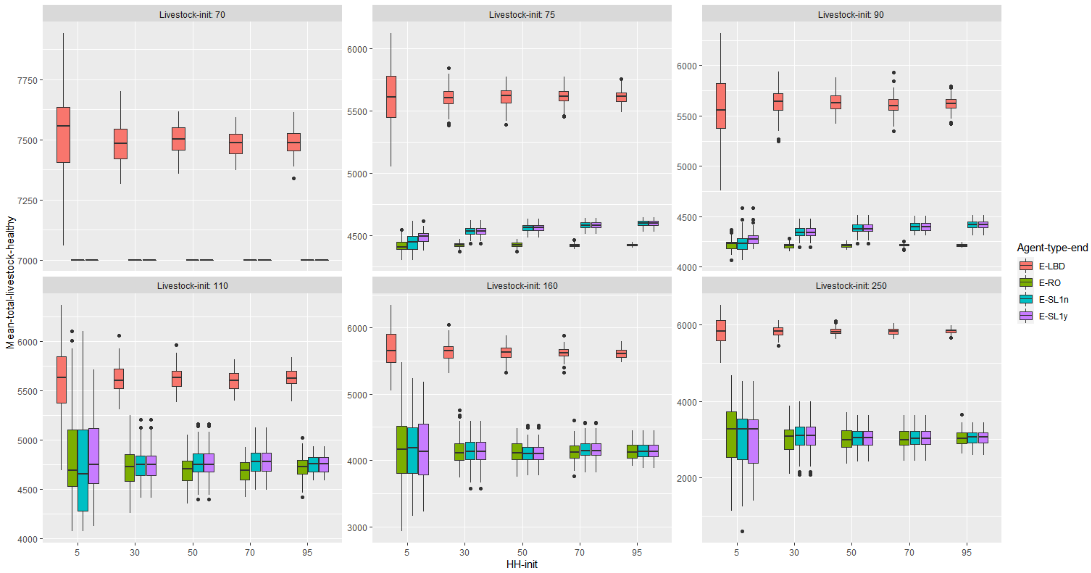

Mean-reserve-biomass-occupied-patches:

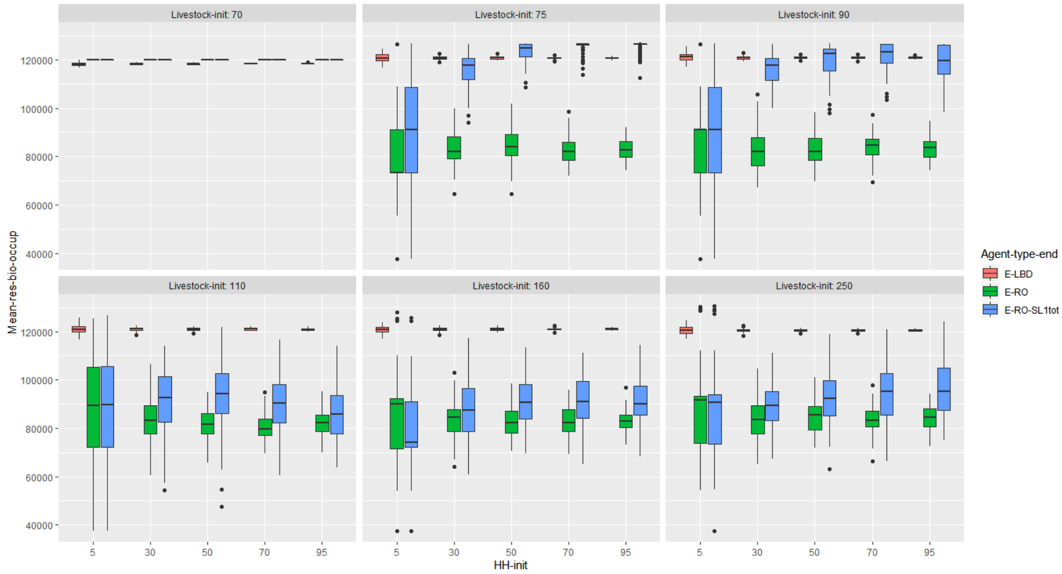

Mean-reserve-biomass-occupied-patches with SL1 results split by learnt/didn't learn:

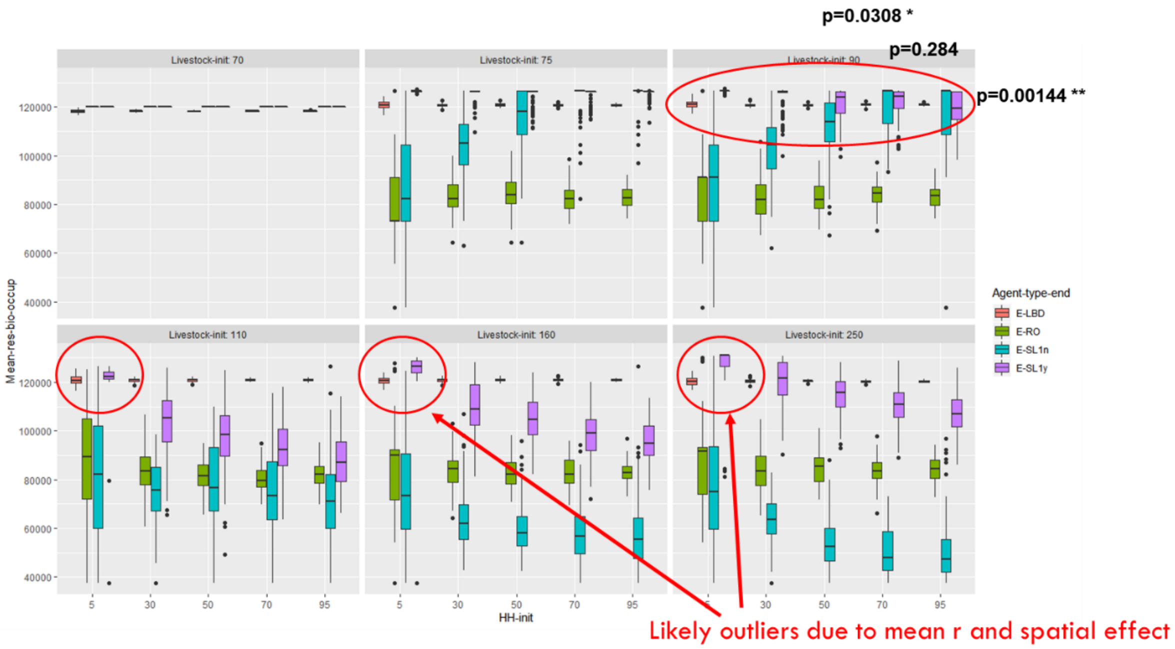

Gini-coefficient

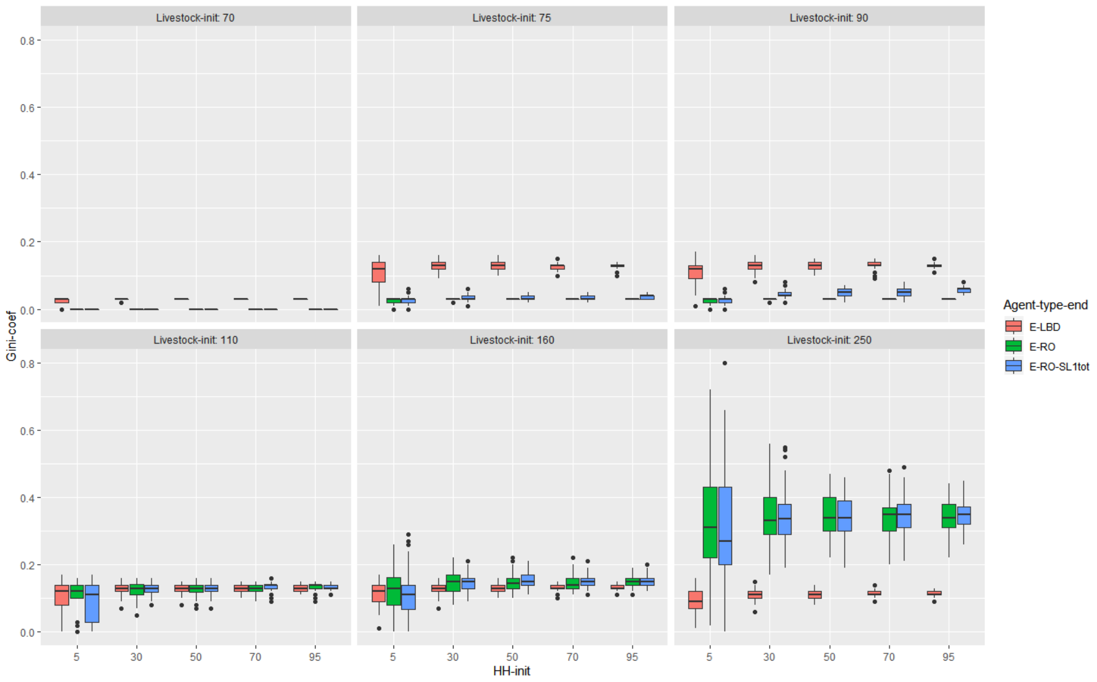

6. Sensitivity analyses knowledge radius, 100 repetitions, HH-init: 5 50 95

Economic variables:

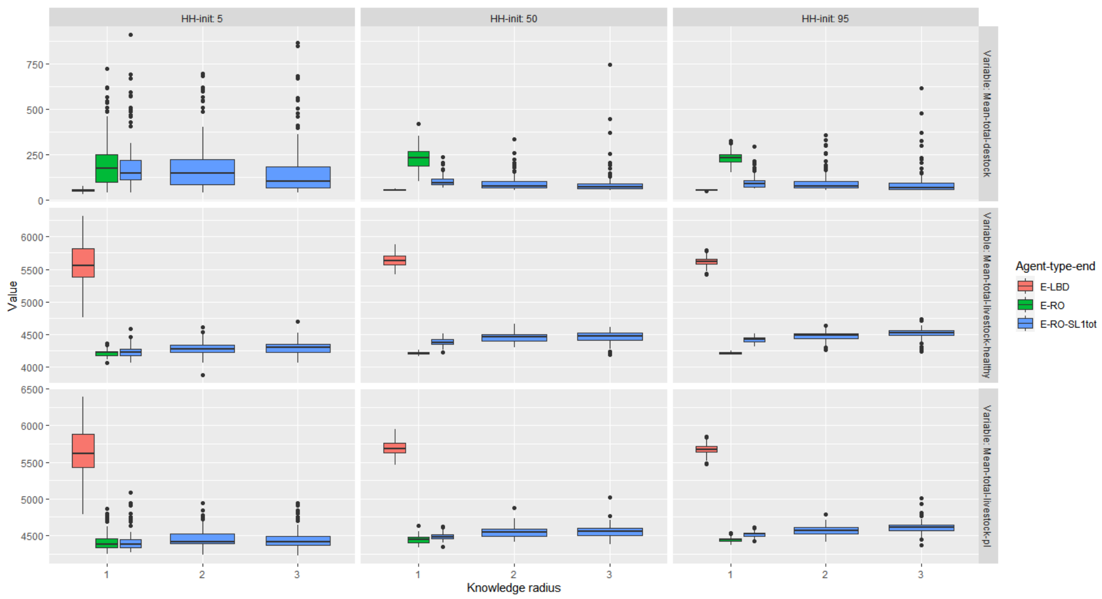

Ecological variables:

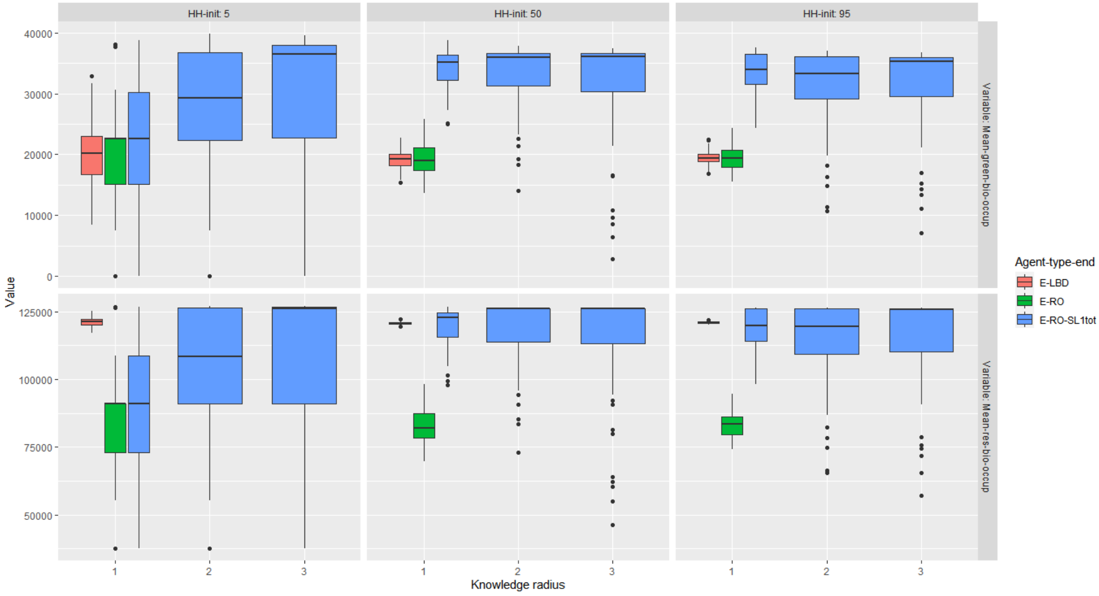

Gini-coefficient: