Electronic Supplementary Material for:

"Improving the representation of smallholder farmers' adaptive behaviour in agent-based models: Learning-by-doing and social learning."

Supplementary Information

Improving the Representation of Smallholder Farmers' Adaptive Behaviour in Agent-Based Models: Learning-by-Doing and Social Learning

C.I. Apetrei, N. Strelkovskii, N. Khabarov, V. Javalera Rincón ⬧ *Ecological Modelling* ⬧ Dec 2023

A. ODD+D Protocol

Extending the RAGE model (Dressler et al., 2018) with learning agents. This ODD+D protocol (Müller et al., 2013; Grimm et al., 2010) should be read together with the one accompanying the original model. It follows the same layout for easier reading. We do not describe the whole model, but focus on the extension module and how it differs from the original.

! N.B. Users who want to run the original model from our version (by disabling the extension module) should note that the initial model dynamics have not been exhaustively verified after our additions.

Outline		Guiding questions	Description
I) Overview	I.i. Purpose	I.i.a What is the **purpose** of the study?	The purpose of the study is to advance the understanding of farmers' decision-making by explicitly accounting for adaptive behaviour and learning, in particular learning-by-doing and social learning. We aim to: a) make a methodological contribution about how learning-by-doing and social learning, two concepts from the adaptive (co)-management literature, might be represented in a social-ecological agent-based model; b) elucidate how these processes and their interactions might affect decision outcomes.
		I.ii.b For whom is the model designed?	Scientists interested to understand and include adaptive behaviour in models of farmers' decision-making.
	I.ii Entities, state variables, and scales	I.ii.a What kinds of **entities** are in the model?	Just like in the original model: - households that make individual decisions about their livestock; - patches that represent pastures for livestock to graze.

		I.ii.b By what **attributes** (i.e. state variables and parameters) are these entities characterized?	In addition to the original model, the extension adds the following attributes: **Households** – attributes added which are important conceptually: - **household-reserve-biomass-memory:** stores the amounts of reserve biomass observed in previous time periods (ticks); defined as array type; - **household-livestock-placed-memory**: stores the number of livestock placed on the pasture in previous time periods; also array; - **household-risk-att**: the so-called “r-parameter”, a random number with a value between -0.95 and 0.95 in increments of 0.05. This is used to represent agent heterogeneity in dealing with uncertainty and it is the degree by which a final decision by the household will deviate from a “rationalized” calculation of the total number of livestock to be placed on the pasture; this value can change during the simulation if the agent is engaged in social learning. **Households**’ states are also captured by the following new variables, which are employed to support data collection and analysis: - **household-risk-att-init**: the random value of the “r-parameter” allocated to the agent at the beginning of the simulation; - **household-calc-diff**: stores latest calculation of how much the biomass on the pasture has changed since previous time period (current observation minus last past observation); - **household-behavioral-type-init**: stores the initial behavioural type of the agent (because the variable in the original model “household-behavioral-type” can change its value during the simulation); - **household-SL?** boolean variable storing whether an agent with behavioural type “social learning” – E-RO-SL1 – has engaged in social learning at least once during the simulation (false: no social learning interaction has taken place; true: at least one social learning interaction has taken place); - **household-SL2?** boolean variable storing whether an agent of any behavioural type has changed its behavioural type (when “strategy switching” ON); - **new-household-behavioral-type:** under “strategy switching”, this variable is used to store the behaviour type of the most successful neighbour, which the current agent will adopt at the beginning of the following time period; - **new-household-risk-att:** under “social learning 1”, this variable is used to store the value of the r-parameter of a most successful neighbour, which the current agent will adopt at the beginning of the following time period;

			- **new-color:** used for changing color in the following time period according to the new-household-behavioural-type; - **livestock-healthy-total:** used to count the total number of healthy livestock that an agent has kept over time - **destock-total:** used to count the total number of livestock that were underfed over time (overshooting the carrying capacity) **Patches**: no changes relative to the original model
		I.ii.c What are the **exogenous factors / drivers** of the model?	- In the extension model, the vegetation regeneration is **no longer** driven by stochastic rainfall; a fixed value is used for rain to enable comparability of pasture conditions. However, the possibility of stochastic rainfall still exists in the model and it is possible to enable it. - The extension model analyses do not employ *resting* policies from the original model; - The **r-parameter (household-risk-att)** is initialized at the beginning of the simulation with a random value (stored in **household-risk-att-init**), but it will change during the run (see also II.ix.a - "Stochasticity")
		I.ii.d If applicable, how is **space** included in the model?	No changes.
		I.ii.e What are the **temporal** and **spatial** resolutions and extents of the model?	No changes.

	I.iii Process overview and scheduling	I.iii.a What **entity** does what, and in what order?	If extension model is active (**extension-model?** is TRUE)
			- **Initialization:** set up households with different **household-risk-att** each, set up patches as in original model, set up fixed rain value for all patches (282.71, value from rain value file of original model)
			- In each **tick**:
			    o **Patches** update their green biomass (with fixed rain)
			    o All **households** (in sequence):
			        ▪ If applicable, households with SL1 agent type (**household-behavioral-type**="E-RO-SL1") **update their r-parameter** with a value retained at the end of the previous tick;
			        ▪ If applicable and if strategy switching is enabled (**SL2-strategy-switch?** TRUE), all households **update their agent type** / household behavioural strategy with the value retained at the end of the previous tick;
			        ▪ Agents **make decisions** on the number of livestock to be placed in the current round, according to their agent type;
			        ▪ Livestock **feeds** on the pasture biomass and values for the destock variable are updated, if not enough biomass
			    o All **households** (in sequence):
			        ▪ If engaged in social learning (either because global strategy switching is enabled, or because agents exist of type E-RO-SL1) agents **observe** the performance of **their neighbours** (in sequence)
			- NB: this step is in a separate "ask households" procedure because all households need to have finished their decisions and calculations before social learning agents start evaluating the performance of their neighbours
			    o **Patches** update their reserve biomass
			    o **Households** add new reserve biomass to their memory
			
			*Livestock reproduction is deactivated in the extended model, as it is not matching the theoretical assumptions of the study accompanying the model and, hence, not used. Modellers interested to further tweak the model and use livestock reproduction with the extension module enabled should note that the livestock growth step comes at **the end** of the **to go** procedure (not at the beginning as in the original model). This is because the variables livestock and destock are being used in the behaviour of extension model agent types.

II) Design concepts	II.i Theoretical and Empirical Background	II.i.a Which general **concepts**, theories or **hypotheses** are underlying the model's design at the system level or at the level(s) of the submodel(s) (apart from the decision model)? What is the link to complexity and the purpose of the model?	We implement three types of learning processes, the first two as agent types, the third as a global behaviour that can be enabled from the visual interface: - learning-by-doing is individual learning based on observations of how the environment reacts to past decisions. This is embedded in the agent type E-LBD. Details on the algorithm are in the procedure **decide-livestock-placed**, in the if-block with **household-behavioral-type = "E-LBD"**; - social learning 1 models agents who imitate the **decision dispositions** of their most successful neighbours (the agent attribute "household-risk-att"). This is embedded in the agent type E-RO-SL1. Basic details on the algorithm are in the procedure **decide-livestock-placed**, in the if-block with **household-behavioral-type = "E-RO" or "E-RO-SL1"**; in addition, this learning type relies on neighbour imitation, so relevant procedures are also: **observe-neighbours** and **update-r-param**; - social learning 2 is a type of double-loop learning and it models agents who imitate the **decision rules** (agent types) of their most successful neighbours. This behaviour is implemented as a global behaviour of strategy switching. Agents of any type may engage in strategy switching concurrently with the learning process defined by their type. Furthermore, the extension model: - no longer assumes a common property resource system; each occupied patch is **the private property** of the household occupying it who graze their herd on it; - households decide every round how many livestock to sell/buy and an unlimited budget as well as supply/demand on an exogenous market is assumed; - livestock for which there is insufficient fodder in one round (**destock**) do not die, they are considered underfed – agents observe the destock value and take it into account in their next decision on herd size.
		II.i.b On what **assumptions** is/are the agents' decision model(s) based?	Decision assumptions are documented in detail in the accompanying paper and mapped to a conceptual framework for modelling learning proposed by the authors. The implementation of learning-by-doing and social learning is based on conceptualisations from the adaptive management literature.
		II.i.c Why is a/are **certain decision model**(s) chosen?	We wanted to compare the outcomes of two different types of learning which are often mentioned in the adaptive management literature: learning-by-doing and social learning. Our model also serves to illustrate the challenges encountered and specifications required to implement learning processes in farmers' ABMs.

		II.i.d If the model / a submodel (e.g. the decision model) is based on **empirical data**, where does the data come from?	Not applicable.
		II.i.e At which level of **aggregation** were the **data** available?	Not applicable.
	II.ii Individual Decision Making	II.ii.a What are the **subjects** and **objects** of decisionmaking? On which level of **aggregation** is decision-making modeled? Are multiple **levels** of decision making included?	- Individual households are the **subjects** of decision making, as well as of learning. - The size of the herd in each round, i.e. the number of livestock to be placed on the pasture, is the **object of decision-making**. - The **object of learning** (what changes as a result of learning) differs for each type of learning studied: o For learning-by-doing: agent actions change (decisions in next round); o For social learning 1: agent attributes change (r-parameter); o For social learning 2: agent learning type changes (behavioural strategy, i.e. agent type).
		II.ii.b What is the basic **rationality** behind agents' decision-making in the model? Do agents pursue an explicit **objective** or have other success criteria?	The implicit goal of the agents is to maximize the number of healthy livestock they can keep at each moment (i.e. to approximate the carrying capacity).

		II.ii.c How do agents make their **decisions**?	When engaged in learning-by-doing, agents make their decisions of how many livestock to place on the pasture in each round based on: a) the sensed available reserve biomass; b) their previous decision on how much livestock to buy/sell; d) the number of livestock that were observed to be underfed in the previous round (destock value). Social learning agents observe their neighbours and use profit maximisation to decide which neighbour to imitate (i.e. they evaluate performance by comparing the numbers of healthy livestock of their neighbours with that of themselves, in the current round). Alternative agent goals could be implemented in following versions of the model.
		II.ii.d Do the agents **adapt** their behavior to changing endogenous and exogenous state variables? And if yes, how?	Yes. All learning processes described here are forms of adaptation. In response to feedbacks from their social-ecological environments, the agents adapt either their actions, their attributes or their behavioural strategy.
		II.ii.e Do **social norms** or cultural values play a role in the decision-making process?	Yes, but only insofar as social learning by imitation is interpreted as norm adoption. However, in this model we do not explicitly deal with institutions as broader societal norms. An institutional layer exists in the original model and could be activated in the extension, if needed.
		II.ii.f Do **spatial aspects** play a role in the decision process?	Yes. Social learning (1 and 2) agents observe 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 so that k=1 delimits the Moore neighbourhood of an agent.

		II.ii.g Do **temporal aspects** play a role in the decision process?	Yes. Learning-by-doing agents consider previously observed pasture states when making their decisions on herd size. In a heterogeneous population with mixed learning-by-doing and social learning 1 agents, and with strategy switching enabled, there will be differences in the pasture state of an agent who switch to social learning 1 agent type only in round N vs. the pasture of an agent who had this behavioural type from the beginning. Thus, there is a co-evolution between the pasture state and the agent's behavioural strategy. Social learning (1 and 2) can only start from the second tick, when the results of the first decision round become available. In tick 2 agents observe the performance of their neighbours and then from tick 3 they can imitate either the r-parameter or the agent type of the most successful neighbour observed.
		II.ii.h To which extent and how is **uncertainty** included in the agents' decision rules?	All agents deviate from the "rationalized" decision which follows from applying the rules of their agent type (learning process) by a certain percentage given by the *r-parameter*. This is meant to introduce agent heterogeneity in dealing with uncertainties. Agent *r-parameter* values are initialized with values drawn from a discrete uniform distribution in increments of 0.05 between -0.95 and +0.95. The "rationalized" decision is the one corresponding to r=0. Agents with r≠0, will decide to place more (if r>0) or less livestock (if r<0) than they would in a world of perfect information. The choice to limit the range of the *r-parameter* to -0.95 and +0.95 instead of -1/+1 was made to avoid an agent giving up the entire livestock herd and removing themselves from the farming activity. **Please note:** the risk-mode? switch in the interface controls the activation/deactivation of agent heterogeneity in dealing with uncertainty (whether the r-parameter values are taken into account in decision-making or not). This switch is only implemented for learning-by-doing agents. Baseline and social learning 1 agents always take r values into account. The risk-mode switch is ON by default. An OFF value is only useful for specific experiments with homogeneous populations of learning-by-doing agents.

	II.iii Learning	II.iii.a Is individual **learning** included in the decision process? How do individuals change their decision rules over time as consequence of their experience?	In a very strict modelling sense, the only **learning** implemented is social learning 2, i.e. strategy switching / double-loop learning, where agents alter their decision **rules.** See details above (II.i.a and II.ii.a). However, we take a soft understanding of the concept of learning, as resulting from adaptive management and sustainability science literatures, to illustrate how "learning-by-doing" and "social learning" could be implemented in an ABM. In modelling terminology, learning-by-doing and social learning 1 would be considered adaptation processes.
		II.iii.b Is **collective learning** implemented in the model?	We only deal with **individual learning** in the model. However, we assess emergent outcomes of strategy switching at the collective level, by looking at agent type diffusion in heterogeneous populations of non-learning, learning-by-doing and social learning 1 agents.
	II.iv Individual Sensing	II.iv.a What **endogenous** and **exogenous** state variables are individuals assumed to sense and consider in their decisions? Is the **sensing process** erroneous?	- Households sense the vegetation state of their own pastures - Households also sense the agent attributes of their neighbours within a *knowledge radius k*, including their total numbers of livestock, as well as the number of **livestock** that were underfed (**destock**). - Sensing is not erroneous, but real-life uncertainties are approximated through the **r-parameter** which introduces agent-specific variations in the final decision (see II.ii.h).
		II.iv.b What **state variables** of which other individuals can an individual **perceive**? Is the sensing process erroneous?	Households are able to sense state variables of other households – r-param, agent type, livestock and destock. See above.

		II.iv.c What is the spatial **scale** of **sensing**?	- Sensing pasture state: own pasture; - Sensing other agents' state: neighbourhood of knowledge radius k.
		II.iv.d Are the mechanisms by which agents obtain **information** modeled explicitly, or are individuals simply assumed to know these variables?	- Gathering information is not modelled explicitly. Individuals are assumed to know the variables mentioned above.
		II.iv.e Are **costs** for cognition and costs for gathering information included in the model?	Not applicable.
	II.v Individual Prediction	II.v.a Which data uses the agent to predict future conditions?	The agent does not predict future conditions, learning-by-doing is reactive. Agents look at information from the past and the present and decide what to do next.
		II.v.b What internal models are agents assumed to use to estimate future conditions or consequences of their decisions?	Not applicable.

		II.v.c Might agents be erroneous in the prediction process, and how is it implemented?	Not applicable.
	II.vi Interaction	II.vi.a Are interactions among agents and entities assumed as **direct** or **indirect**?	Interactions are indirect and only for social learning 1 agents or when strategy switching is enabled (social learning 2). Agents who are strictly learning-by-doing (and strategy switching is not enabled) do not interact with others, although others might imitate them.
		II.vi.b On what do the interactions depend?	On learning types and on whether strategy switching is activated. Interactions also depend on whether agents are in each other's k radius. On a sparsely populated map (with no or few neighbourhoods overlapping), most agents will not interact.
		II.vi.c If the interactions involve communication, how are such communications represented?	Not applicable.
		II.vi.d If a coordination network exists, how does it affect the agent behaviour? Is the structure of the network imposed or emergent?	Not applicable.

	II.vii Collectives	II.vii.a Do the individuals form or belong to aggregations that affect, and are affected by, the individuals? Are these aggregations imposed by the modeller or do they emerge during the simulation?	Not applicable.
		II.vii.b How are collectives represented?	Not applicable.
	II.viii **Heterogeneity**	II.viii.a Are the agents heterogeneous? If yes, which state variables and/or processes differ between the agents?	Yes, agents are heterogeneous in dealing with uncertainty, as represented by the r-parameter (**household-risk-att**). If the model is run with mixed populations, then agents also differ in agent type according, as per initialization numbers and, if applicable, behaviour diffusion (when strategy switching is on).
		II.viii.b Are the agents heterogeneous in their decisionmaking? If yes, which decision models or decision objects differ between the agents?	Yes, see discussion about learning processes above.

	II.ix Stochasticity	II.ix.a What processes (including initialization) are modeled by assuming they are random or partly random?	There are three sources of stochasticity realized at the initialization of each model run: 1) the distribution of the *r*-parameter in the population, 2) the spatial distribution of agents on the map, and 3) in runs with heterogeneous populations: the allocation of behavioural types to agents. In addition, inherent to NetLogo, households execute actions in a random order.

	II.x Observation	II.x.a What data are collected from the ABM for testing, understanding, and analyzing it, and how and when are they collected?	In addition to the original model, we collect the following data: In the <u>graphical user interface</u>, we plot at each timestep: - Mean reserve biomass of occupied patches by type of agent; - Mean livestock placed on the pasture by type of agent; - The current livestock placed by each household; - The livestock placed on the pasture by Households 1-5 (this gives an indication of how five random agents behave); - For each agent type, we also display: the mean total livestock and the mean total destock (cumulative over all timesteps); - The Gini-coefficient (of total livestock healthy) is also displayed. For the <u>experiments in Behavior Space</u>, we collect: - Mean-R-parameter of the entire population: at the beginning and at the end of the simulation; - The list of all R-parameters of the entire population: at the beginning and at the end of the simulation; - The mean reserve and green biomass of all patches, as well as of occupied patches only, grouped by agent type; - The mean livestock-placed on the pasture in the last timestep (across all agents); - The mean destock in the last timestep (across all agents); - The number of households which changed their r-parameter as a result of social learning 1 (household-SL?=true); - The mean **total livestock placed** and **total destock** (total = cumulative for all timesteps) of agents, grouped by agent type; - The mean total livestock placed and total destock of agents who changed their r-parameter as a result of social learning 1; - Gini-coef (corresponding to procedure in original model) = Gini-coef-surviving: the Gini index based on the values of livestock-healthy of all agents at the end of the last timestep; - Gini-coef-total-livestock-healthy: the Gini index relatively to the values of total-livestock-healthy (cumulative over entire simulation) at the end of the simulation. In addition, for <u>diffusion experiments in Behaviour Space</u>, we collect: - The number of households of each type (at the end of the simulation)

		- The number of households that changed their r-parameter at least once (household-SL?=TRUE) - The number of households that changed their behavioural strategy at least once (household-SL2?=true) Not all these variables are used in all the experiments discussed in the accompanying manuscript. For details, see the manuscript's Supplementary Information D.
		II.x.b What key results, outputs or characteristics of the model are emerging from the individuals? (Emergence)	In all simulations we can observe the effects of decision-making over the state of individual pastures, households' economic performance (number of livestock) and economic inequality in the entire population. In simulations where agents interact (due to social learning 1 and/or 2), we can also observe the diffusion of specific r-parameter values or of behavioural types.
	III.i Implementation Details	**III.i.a How has the model been implemented?**	The model has been implemented in NetLogo 6.1.1.
		III.i.b Is the model accessible and if so where?	The model is available at ComSES Net: "Learning Extension - RAGE RAngeland Grazing Model (version 1.0.0)". https://www.comses.net/codebases/e103feef-2785-41e6-affb-8306c979e83c/releases/1.0.0/
	III.ii Initialization	III.ii.a What is the initial state of the model world, i.e. at time t=0 of a simulation run?	- Like in the original model, all households are initialized with the same number of livestock, same pasture conditions and same fodder requirements. - Households are initialized with different values for the r-parameter (they are heterogeneous in how they deal with uncertainty)

		III.ii.b Is initialization always the same, or is it allowed to vary among simulations?	- Initialization varies randomly in terms of:
			  ○ Distribution of r-parameter values
			  ○ Spatial distribution of agents
			  ○ Which agents are allocated to which behavioural type

		III.ii.c Are the initial values chosen arbitrarily or based on data?	- The number of initial households and the number of initial livestock are input parameters selected based on sensitivity analyses - Other input parameters (e.g. for the vegetation submodel) are the same as in the original model - R-parameter values are generated randomly from a discrete uniform distribution (see II.ii.h).
	III.iii Input Data	III.iii.a Does the model use input from external sources such as data files or other models to represent processes that change over time?	No external data is used.

	III.iv Submodels	III.iv.a What, in detail, are the submodels that represent the processes listed in 'Process overview and scheduling'?	The following submodels run when the extension module is enabled:
			
			Initialization:
			Procedures: *setup, setup-household-behavioral-types*
			- **The extension model is designed to run a maximum of timesteps**. This is because agent memory is initialized as a vector of length 501. For longer runs, memory needs to be extended first;
			- All patches are initialized with the same amounts of green and reserve biomass, and according to the settings in the original model;
			- A fixed rain value is set for all patches;
			- N_H households are created and initialized with the selected behavioural types;
			    o **Please note: in the current implementation it is not possible to mix agent types from the original model with agent types from the extension model.**
			- Memory vectors and r-parameter values are initialized for all households;
			- Household shapes are set up depending on the r-parameter value: triangle for positive r-parameter values, and upside-down triangle (triangle-down) for negative r-parameter values.
			
			Vegetation growth: same as original model
			
			Updating r-parameter of all households of E-RO-SL1 type:
			Procedure: *update-r-param*
			- This procedure can only be triggered from tick 3 onward (see II.ii.g);
			- **If** in the previous round a relevant observation of a neighbour to be imitated has been made (procedure **observe-neighbours**), the r-parameter value to be copied is already stored in the variable **new-household-risk-att**. **Then:**
			    o household-risk-att is updated to the new value;
			    o the colour of the household flashes in the interface;
			    o the shape of the household is adjusted (triangle or triangle-down)
			    o the agent variable indicating that social learning type 1 has taken place at least once (SL-learning?) changes to TRUE.
			
			Updating agent type (behavioural strategy) of all households, if strategy switching is on
			Procedure: *update-household-strategy*
			- This procedure can only be triggered from tick 3 onward (see II.ii.g);

			- **If** in the previous round an observation of a neighbour to be imitated has been made (procedure **observe-neighbours**), the agent type to be copied is already stored in the variable **new-household-behavioral-type. Then:** o household-behavioral-type is updated to the new value; o the colour of the household flashes in the interface; o the agent variable indicating that social learning type 2 has taken place at least once (SL2-learning?) changes to TRUE. **Household decision-making** Procedure: *decide-livestock-placed* (note that for agent types in the original model, a different procedure is used for household decision-making) - For learning-by-doing agents (E-LBD): o **Learning can only happen from the second tick onward** (as pasture condition after first decision needs to be evaluated); o Agent observes the difference in reserve biomass compared to previous round; o The rate of change of the pasture is calculated; o Agent calculates by how much they should adjust their herd size (livestock-buy-sell) compared to previous round in a way that is proportional to the observed rate of change of the pasture's reserve biomass. This value is then compared to the number of livestock that were observed to be underfed in the previous round (if any). The minimum between these two values will represent the adjustment to the last herd size. ▪ For instance, suppose that the reserve biomass has reduced by 7% and a similar reduction of 7% in herd size corresponds to selling 2 livestock. However, the agent knows that at the end of the last round 3 livestock were underfed. Then they will adjust their decision down by 3 livestock, instead of 2. o Agent calculates final value for how many livestock to sell/buy **(livestock-buy-sell-new)** as a deviation by *r* from the value calculated in the previous step; o Final decision on herd size (how much livestock to be placed on the pasture) is taken ensuring that no negative values are taken (floor to 0), e.g. in limit cases when the initial livestock is very big; o Place value of final decision in agent memory. - For baseline and social learning agents (E-RO, E-RO-SL1):

			o Agents observe how many livestock were underfed in previous round (destock value) o The destock value is adjusted to deviate by the *r-parameter* value. o Final decision on herd size is calculated as: *herd size in the last round* minus *adjusted destock*. For a schematic and formalized representation of agents' decision algorithms, please also refer to the Supplementary Information B of the accompanying manuscript. **Livestock feeding:** slightly adjusted from the original model Procedure: *livestock-feed* - Calculate livestock fodder needed - Reserve-biomass-edible is the maximum reserve biomass that may be consumed (based on the **gr2** – grazing harshness – parameter) – see original model ODD; - If all green biomass and reserve biomass-edible consumed, set **destock** to the number of livestock that had nothing left to consume (underfed); - If reserve biomass edible still available, set **destock** to 0 (none were underfed); - For each agent, calculate cumulative values for total number of livestock that were healthy (**livestock-healthy-total**) and total number of livestock that were underfed (**destock-total**). **Neighbour observation:** Procedure: *observe-neighbors* - **Note: This procedure is run in a separate "ask household" loop to ensure that all households have taken their decisions and the new results are available before making performance comparisons;** - Identify neighbouring patches within knowledge radius k; - Check if any neighbours found and if there are neighbours more successful than self (in the number of healthy livestock in the current round); - Note the r-parameter value and agent type of the most successful neighbour and retain them in the memory, if different from own properties.

		III.iv.b What are the model parameters, their dimensions and reference values?	The following parameter values, mostly related to the vegetation submodel, are used exactly as in the original model:

Parameter name	Value	Description
w	0.8	biomass growth rate
gr1	0.5	grazing pressure on green biomass
gr2	0.1	grazing pressure on reserve biomass
rue	*0.002*	rain use efficiency (rue)
mg	*0.1*	mortality green biomass
mr	*0.05*	mortality of reserve biomass
Rmax	*150000*	maximum reserve biomass
d	*1/Rmax*	density dependence of reserve biomass, d=1/Rmax
R0part	*0.6*	initial reserve biomass as fraction of green biomass
lambda	*0.5*	growth limit of green biomass
intake	*640*	fodder consumed by one livestock unit

In addition, the following global parameter values are used in the extension model:

Parameter name	Value	Description
number-pastures	*100*	total number of pastures
rain	*282.71*	Rain falling on pastures – this is the first value from a fixed list of values provided with the original model
b	*0*	livestock birth rate
risk-mode?	*ON*	Agent heterogeneity in dealing with uncertainty - only implemented for learning-by-doing agents.
ticks	*Default 100* *Max 500*	Number of ticks that the model runs; values higher than 500 need manual adjustment of agent memory size
knowledge-radius	*1*	The size of the neighbourhood in which a social learning agent evaluates neighbours' performance is the Moore neighbourhood. Sensitivity analyses were conducted for higher values.

		III.iv.c How were submodels designed or chosen, and how were they parameterized and then tested?	Decision submodels are following theoretical considerations of learning-by-doing and social learning in the context of smallholder farmers’ adaptive management. They were tested following design-of-experiments principles and comprehensive one-factor-at-a-time sensitivity analyses to determine input values for our parameters.

References

Dressler, G., Groeneveld, J., Buchmann, C.M., Guo, C., Hase, N., Thober, J., Frank, K., Müller, B., 2018. RAGE RAngeland Grazing Model. NetLogo. CoMSES Computational Model Library.

Grimm, V., Berger, U., DeAngelis, D.L., Polhill, J.G., Giske, J., Railsback, S.F., 2010. The ODD protocol: A review and first update. Ecological Modelling 221 (23), 2760–2768.

Lindkvist, E., Norberg, J., 2014. Modeling experiential learning: The challenges posed by threshold dynamics for sustainable renewable resource management. Ecological Economics 104, 107–118.

Müller, B., Bohn, F., Dreßler, G., Groeneveld, J., Klassert, C., Martin, R., Schlüter, M., Schulze, J., Weise, H., Schwarz, N., 2013. Describing human decisions in agent-based models – ODD + D, an extension of the ODD protocol. Environmental Modelling & Software 48, 37–48.

Schlüter, M., Müller, B., Frank, K., 2019. The potential of models and modeling for social-ecological systems research: the reference frame ModSES. E&S 24 (1).

B. 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.

Here, the behaviours are presented in a formalized manner. A simpler interpretation of the following algorithms can be found in the ODD+D protocol, section III.iv.a, where the "Household decision-making" procedure is explained in natural language.

B.1. Baseline agent type E-RO

The behaviour of **baseline agents** is represented in Figure B1.1, where t represents the current time step, $L(t\text{-}1)$ is the number of livestock placed in the previous round, $L_d(t\text{-}1)$ is the number of livestock that were observed to be hungry/underfed 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 the r-parameter $r=0$. 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.:

$$L_R = L(t-1) - L_d(t-1) \tag{1}$$

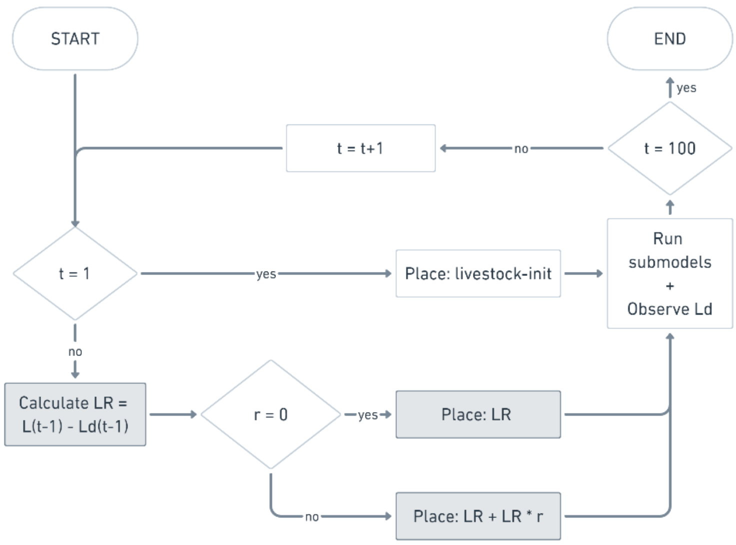

Figure B1.1 Schematic representation of E-RO agent behaviour, i.e. baseline / no learning agents

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

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

B.2. Learning-by-doing type E-LBD

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

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

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:

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

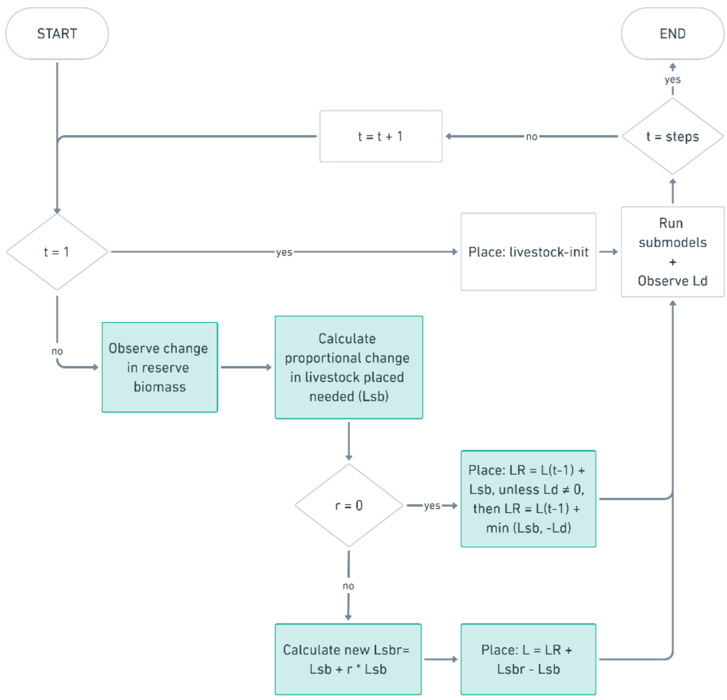

Figure B2.1. Schematic representation of E-LBD agent behaviour

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 ***r≠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:

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

B.3. Social learning type E-SL1

Social learning agents compare their own performance to the performance of their neighbours (from the neighbourhood defined by the knowledge radius *k*). They imitate the *r-parameter* of the most successful neighbour, i.e. the one with the highest number of livestock healthy. The algorithm is shown in Figure B3.1.

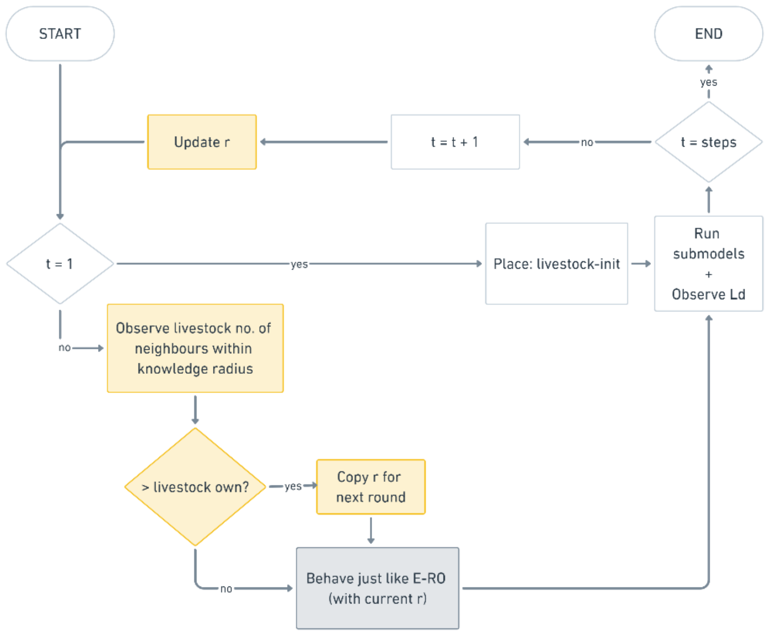

Figure B3.1. Schematic representation of E-SL1 agent behaviour

C. Strategy switching as “Social Learning 2”

A representation of strategy switching is shown in Figure C.1. In the code, this is implemented with two procedures described in detail in the ODD+D protocol, in section III.iv.a: **neighbour observation** and **updating agent type (behavioural strategy)**

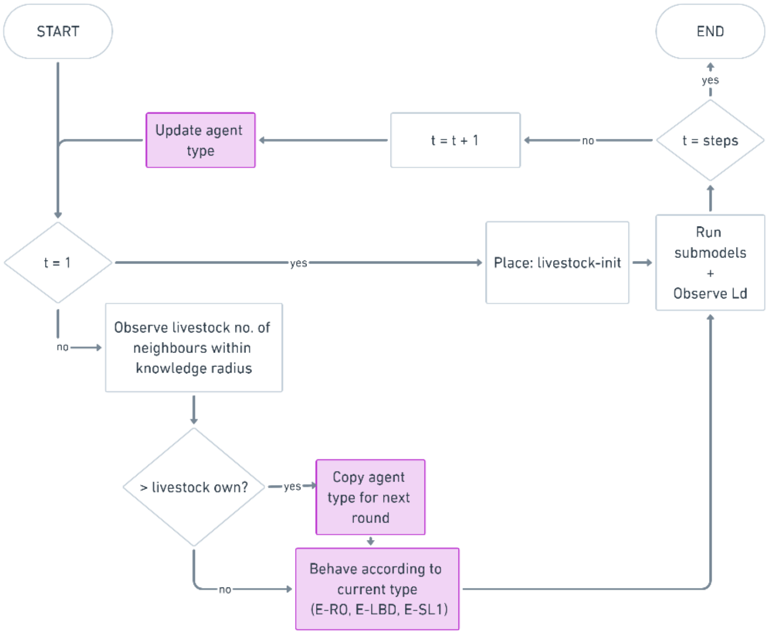

Figure C.1. Schematic representation of Social Learning 2 behaviour: strategy switching

D. Variables for Experiments 1-3

Here we provide additional details on the dependent and control variables in Table 2 of the main manuscript.

Dependent variables. In experiments 1 and 2 we measure economic, ecological and social outcomes, operationalised with three key variables: the *mean total livestock healthy* over the entire period of simulation, the *mean reserve biomass* at the end of the simulation and the *Gini-index of total livestock healthy* over the entire population. For model verification, we also include results on other variables.

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. It is the sum of the livestock sustained at each time step, i.e. the sum of the differences between the livestock placed (L) and the livestock that was underfed (destock, L_d):

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

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 at the end of the simulation. The dependent variable of interest is, thus, the *mean total livestock healthy*, μ_{T_i}:

$$\mu_{T_i} = \frac{\sum_{i=1}^{N} T_i}{N}, \tag{2}$$

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. Unoccupied patches are not included, because they all have the same vegetation dynamics and are not affected by household decisions.

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 healthy at the end of the simulation is distributed across the entire population of agents. A higher Gini-index represents a larger difference in livestock numbers between the most and least endowed farmers.

For experiment 3, the dependent variable is the *number of agents* of each behavioural type at the end of the simulation.

Control variables. These were used for sensitivity analyses to ensure robustness of our final input parameters (see Supplementary Information – Part E).

Population size denotes the total number of agents (households) that the model world is initialised with. *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. The *knowledge radius* is relevant only for social learning behaviour. Model *seed* is a control for stochastic processes. A fixed model seed has been used in sensitivity analyses for simulations with one agent. The reported experiments, on the other hand, were run with a large number of *repetitions* to even out stochastic effects of spatial placement and of *r-parameter* initialization. Because the *r-parameter* of each agent 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 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 without having a fixed seed for the model. Lastly, *time steps* defines the number of decision-rounds after which the model stops and outcomes are compared.

E. Input Parameters for Experiments 1-3

Table E.1. Input parameters for Experiment 1 (total runs: 3000). *Highlighted: parameter corresponding to the independent variable.*

Parameter name	Value	Description
Homog-behav-types?	On (=homogeneous)	Distribution of behavioural types in the agent population
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

Table E.2. Input parameters for Experiment 2 (total runs: 200). *Highlighted: parameter corresponding to the independent variable.*

Parameter name	Value	Description
Homog-behav-types?	Off (=heterogeneous)	Distribution of behavioural types in the agent population
Number-E-RO	20	Initial number of agents with behavior type E-RO
Number-E-LBD	20	Initial number of agents with behavior type E-LBD
Number-E-RO-SL1	20	Initial number of agents with behavior type E-RO-SL1
HH-init (number-households)	60	Population size / initial number of households
Livestock-init	90	Initial herd size
SL2-strategy-switch	*2 factorial levels:* On, Off	Social learning 2 / global strategy switching: active or inactive
Timesteps	100	Simulation length
	Random	Model seed
	1000	Repetitions

Table E.3. Input parameters for Experiment 3 with three sub-experiments (total runs: 900 + 900 + 200 = 2000). *Highlighted: parameters corresponding to the independent variables.*

Sub-experiment	Parameter name	Value	Description
Diffusion LBD *(900 runs)*	**HH-init (number-households)**	*3 factorial levels* {15; 45, 85}	Population size / initial number of households
	Number-E-LBD	*3 factorial levels* {1; 3; 5}	Initial number of agents with behavior type E-LBD
	Number-E-RO = Number E-RO-SL1	*(HH-init – Number-E-LBD)* / 2	Initial number of agents with behavior type E-RO and E-RO-SL1 – equal proportions
Diffusion SL1 (similar to the above) *(900 runs)*	**HH-init (number-households)**	*3 factorial levels* {15; 45, 85}	Population size / initial number of households
	Number-E-RO-SL1	*3 factorial levels* {1; 3; 5}	Initial number of agents with behavior type E-RO-SL1
	Number-E-LBD = Number E-RO	*(HH-init – Number-E-SL1)* / 2	Initial number of agents with behavior type E-LBD and E-RO – equal proportions
"Battle" of learning types *(600 runs)*	**HH-init (number-households)**	*6 factorial levels* {3; 9; 15; 30; 60; 90}	Population size / initial number of households
	Number-E-LBD = Number E-RO = Number E-RO-SL1	*HH-init* / 3	Initial number of agents for each behavior type – equal proportions
All sub-experiments	**Homog-behav-types?**	Off (=heterogeneous)	Distribution of behavioural types in the agent population
	Livestock-init	90	Initial herd size
	SL2-strategy-switch	On	Social learning 2 / global strategy switching: active
	Timesteps	100	Simulation length
		Random	Model seed
		1000	Repetitions

Table E.4. Auxiliary global parameters used in all experiments. *These parameters are not directly related to the research questions but govern the ecological processes described by submodels. The default input values from the original RAGE model were employed.*

Parameter name	Value	Description
w	0.8	biomass growth rate
gr1	0.5	grazing pressure on green biomass
gr2	0.1	grazing pressure on reserve biomass
rue	*0.002*	rain use efficiency (rue)
mg	*0.1*	mortality green biomass
mr	*0.05*	mortality of reserve biomass
Rmax	*150000*	maximum reserve biomass
d	*1/Rmax*	density dependence of reserve biomass, d=1/Rmax
R0part	*0.6*	initial reserve biomass as fraction of green biomass
lambda	*0.5*	growth limit of green biomass
number-pastures	*100*	total number of pastures
rain	*282.71*	Rain falling on pastures – this is the first value from a fixed list of values provided with the original model
b	*0*	livestock birth rate
intake	*640*	fodder consumed by one livestock unit
knowledge-radius	*1*	radius delimiting agents' neighbourhood, here the Moore neighbourhood is used

F. Sensitivity Analyses

Please note: in some of the following figures, the r-parameter value is shown labelled as “Risk” – this is because in our original interpretation *r* represented an agent’s risk attitude. We later extended our interpretation of *r* to include more broadly how agents deal with uncertainty, in line with our literature review. Here we kept the word “Risk” to be consistent with the variable name in NetLogo (household-risk-att).

F.1. One household: ECON outcomes for different livestock-init and r-param

Sensitivity analysis for the effect **of initial livestock numbers** and different values of the **r-parameter** on **economic decisions** (baseline and learning-by-doing agents). **Livestock-placed-end** is the size of the herd in the last decision round. **Destock-end** is the number of underfed livestock in the last round.

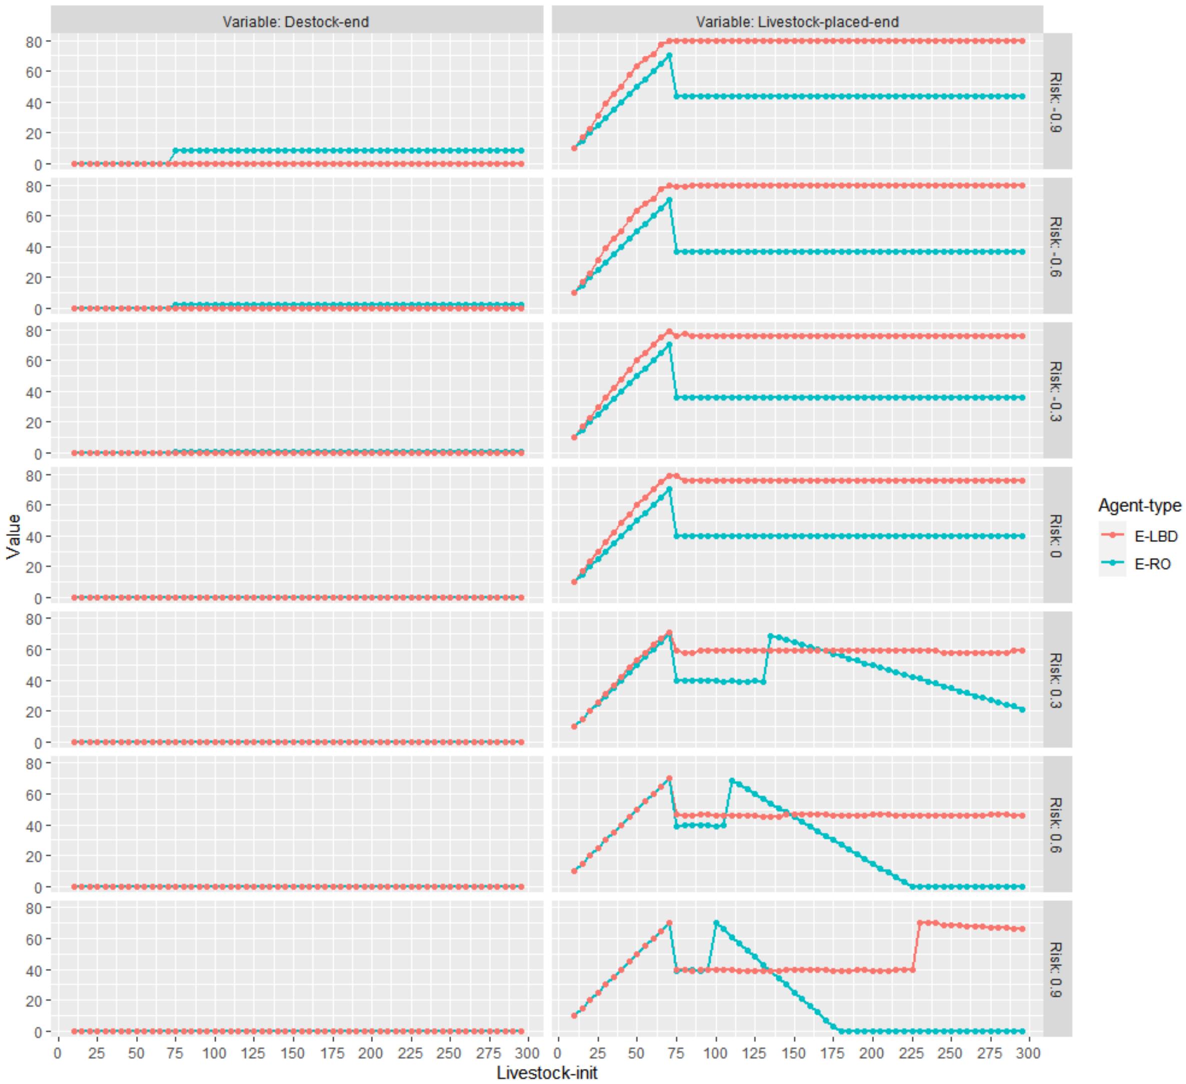

Figure F1.1. Sensitivity analysis for economic effects of the initial herd size and the initial r-parameter values (for one baseline and one learning-by-doing agent, respectively)

F.2. One household: ECOL outcomes for different livestock-init and r-param

Sensitivity analysis for the effect of **initial livestock numbers** and different values of the **r-parameter** on **reserve biomass and green biomass** (baseline and learning-by-doing agents).

a) **Reserve biomass:**

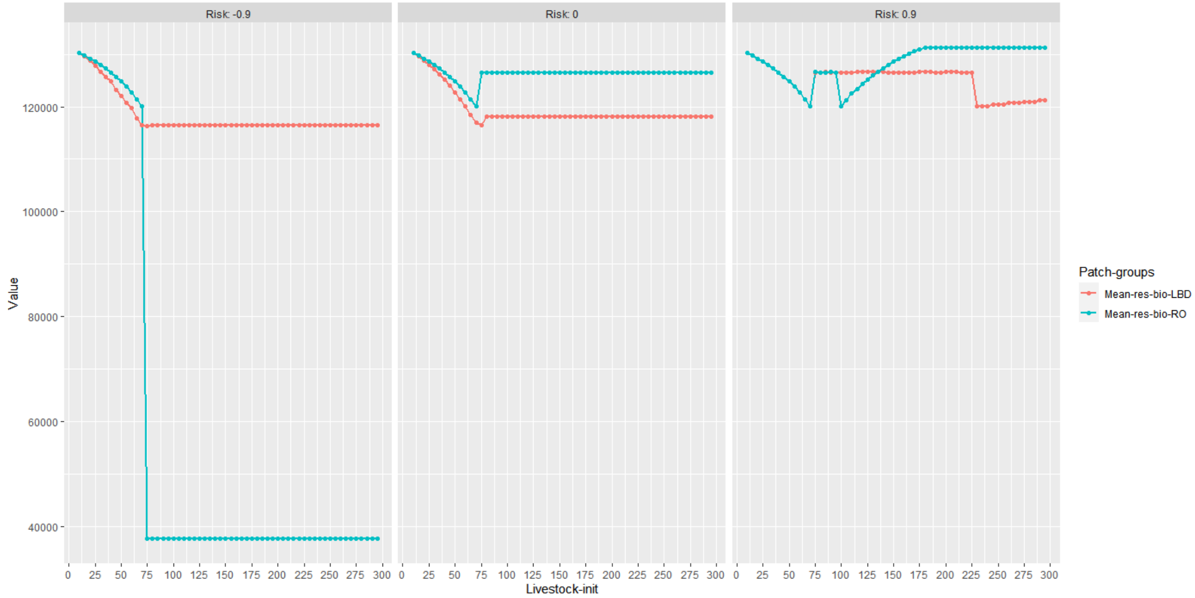

b) **Green biomass:**

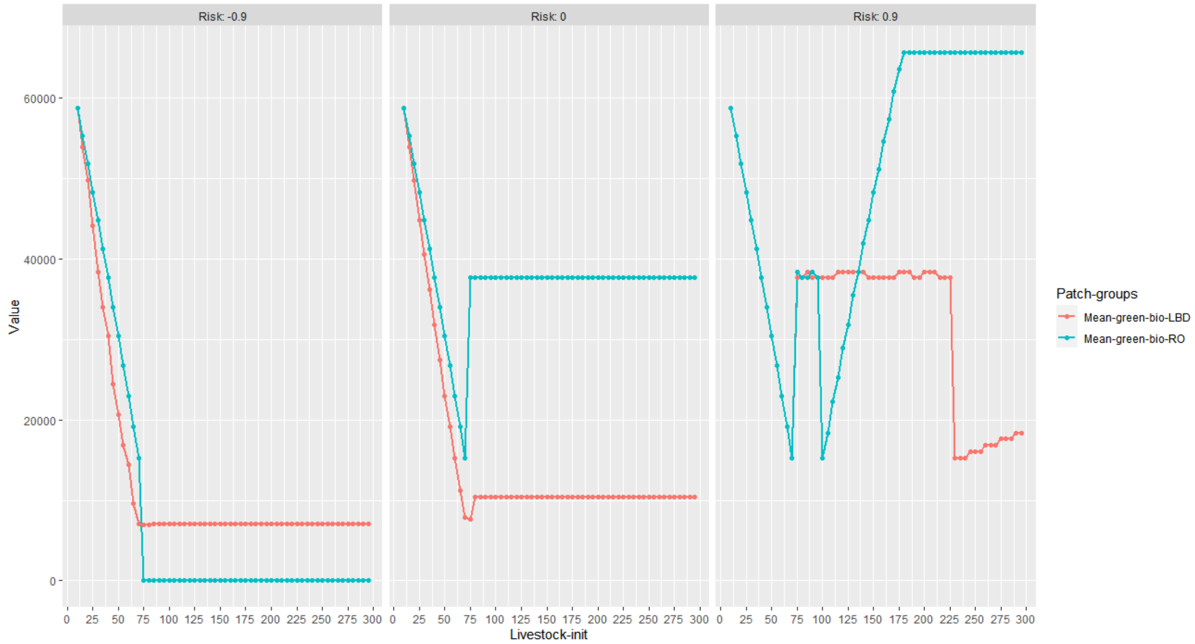

Figure F2.1 Sensitivity analysis for ecological effects of the initial herd size and the initial r-parameter values (for one baseline and one learning-by-doing agent, respectively). Results shown for a) reserve biomass; b) green biomass.

F.3. One household: ECON outcomes for wide range of livestock-init values (up to 1500)

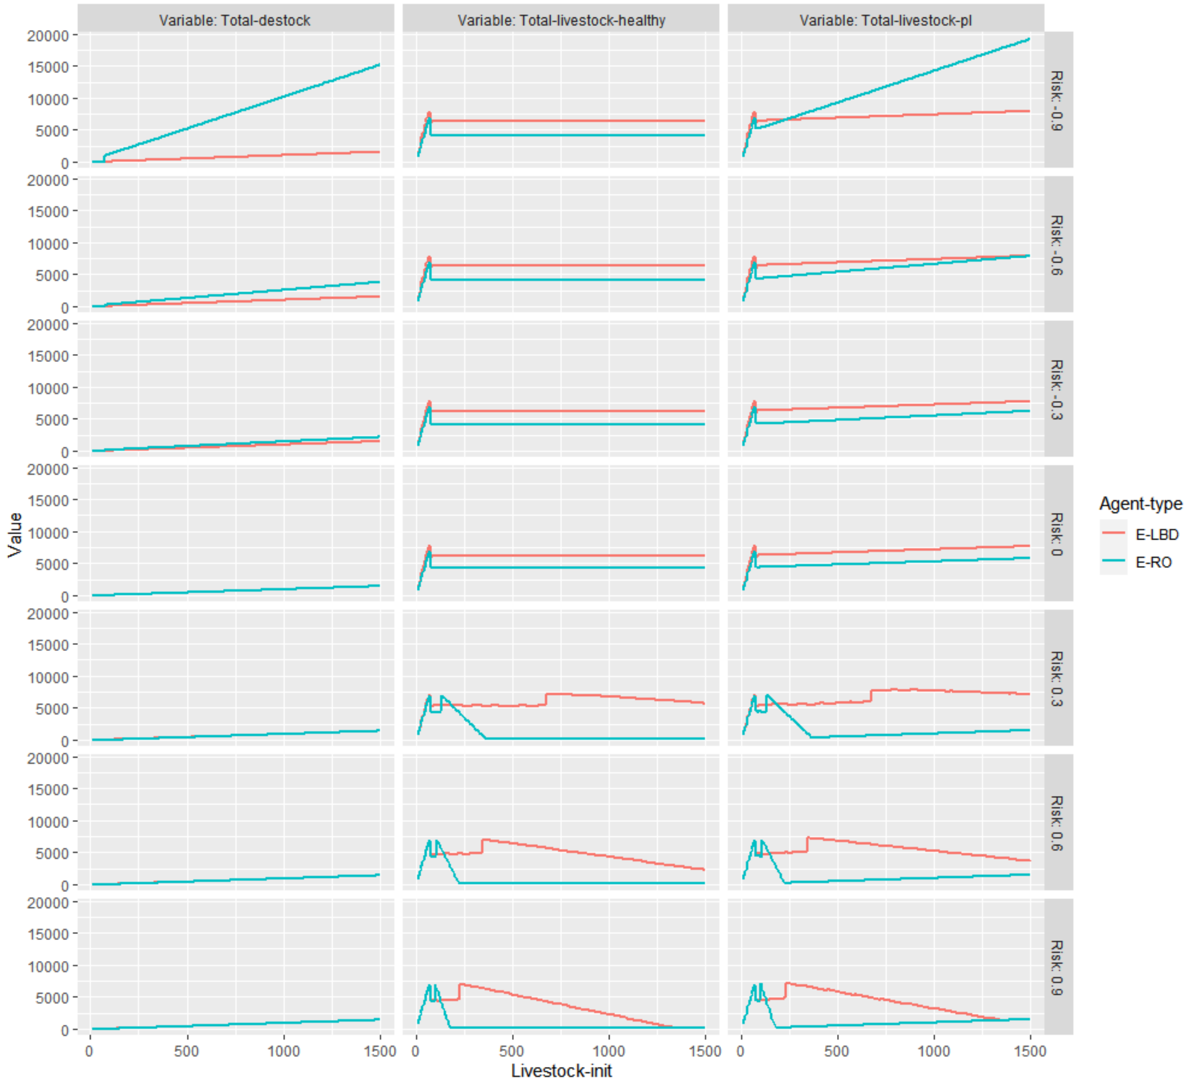

Figure F3.1. Sensitivity analysis for economic effects of VERY HIGH initial herd size and of initial r-parameter values (for one baseline and one learning-by-doing agent, respectively)

F.4. One household: total livestock healthy heat map

This plot explores the combinations of **livestock-init** and **r-parameter** yielding highest numbers (darkest pixels) of total livestock healthy after 100 time steps.

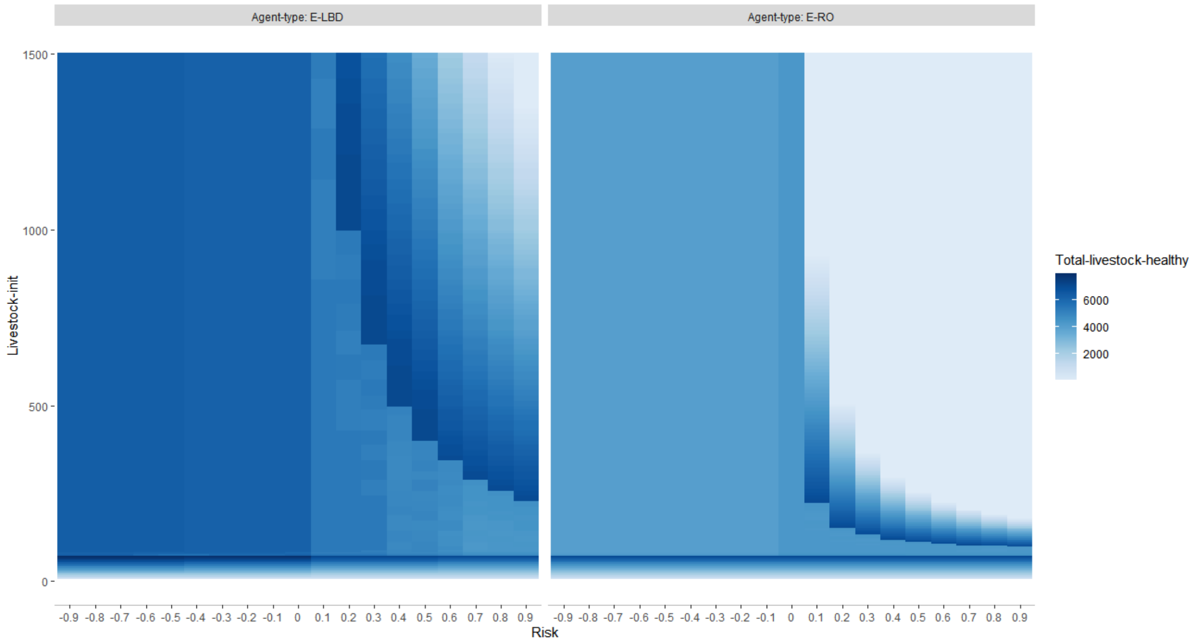

Figure F4.1. Sensitivity analysis for economic effects of the initial herd size (livestock-init) on the mean total-livestock-healthy of agents, as grouped by their r-parameter values.

We also explore the differences in total-livestock-healthy between E-LBD and E-RO agents in the same treatment of livestock-init + r-parameter. In Figure F4.2, blue indicates that E-LBD is performing better, while red indicates that E-RO is performing better (the difference is negative). For readability, negative values are shown as pixels with a dot“.”.

In simulations with many agents, the mean r-parameter value will converge to zero (see Figure F5.1), so the only threshold that we need to account for in further experimental setups is the one at livestock-init=75, across the entire spectrum of r-parameter values. Experimental designs will include treatments with livestock-init values < 75, as well as >75.

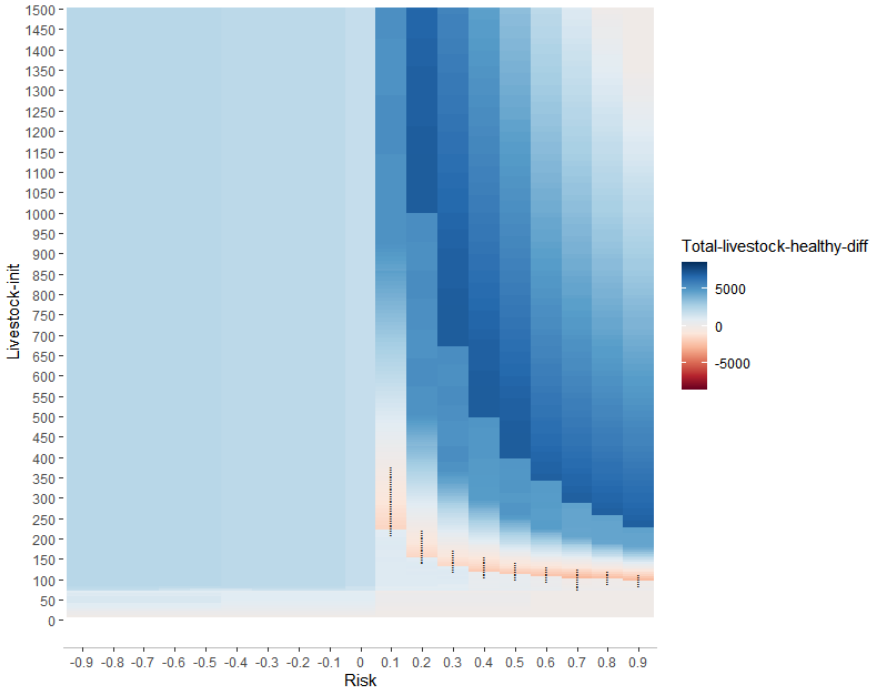

Figure F4.2. Sensitivity analysis exploring differences in economic performance (on total-livestock-healthy) between E-LBD and E-RO agents with the same r-parameter values. Blue values indicate that E-LBD is performing better, red values that E-RO is performing better.

For biomass the results are as following:

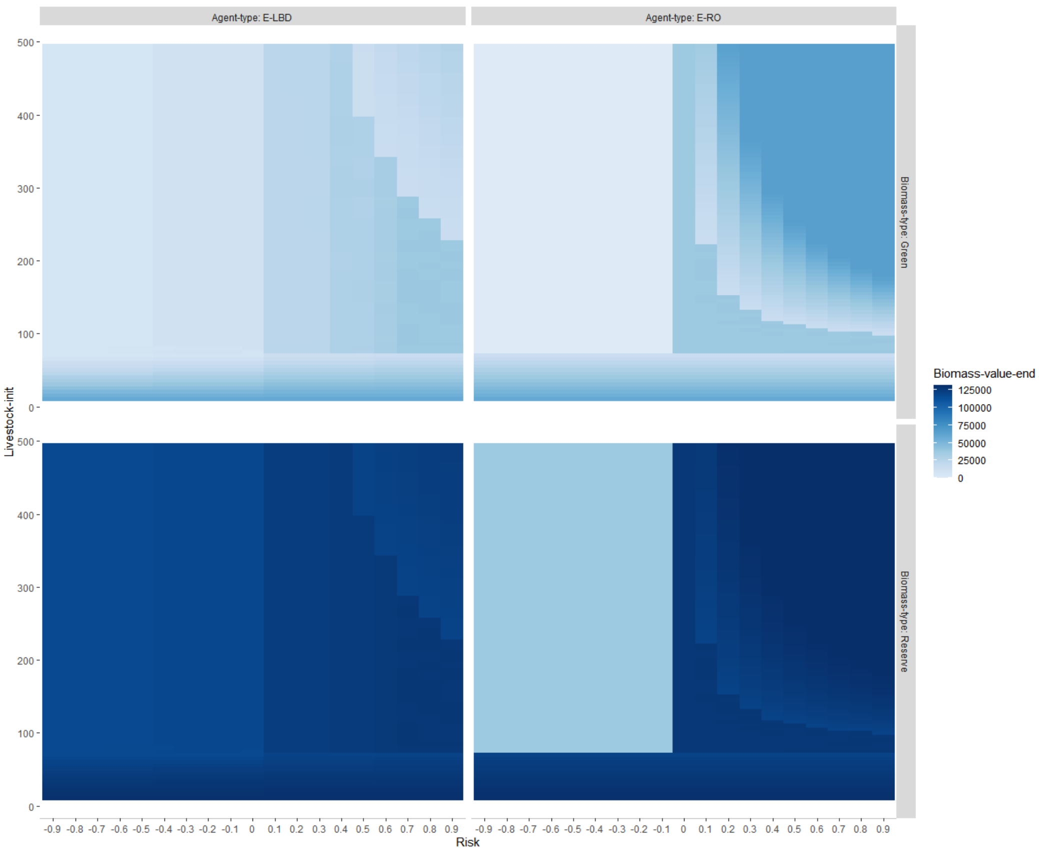

Figure F4.3. Sensitivity analysis for ecological effects of the initial herd size (livestock-init) on the mean reserve and, respectively, green biomass on patches occupied by households with specific r-parameter values.

F.5. R-Mean distribution analysis (1000 repetitions)

The mean values of R-parameter distributions at the beginning of the simulation are shown for 1000 repetitions, different agent types and different initial number of households (HH-init).

- Mean-R (at the beginning of the simulation) converges to 0
- For more than 45 initial households – mean R will be between -0.2 and 0.2

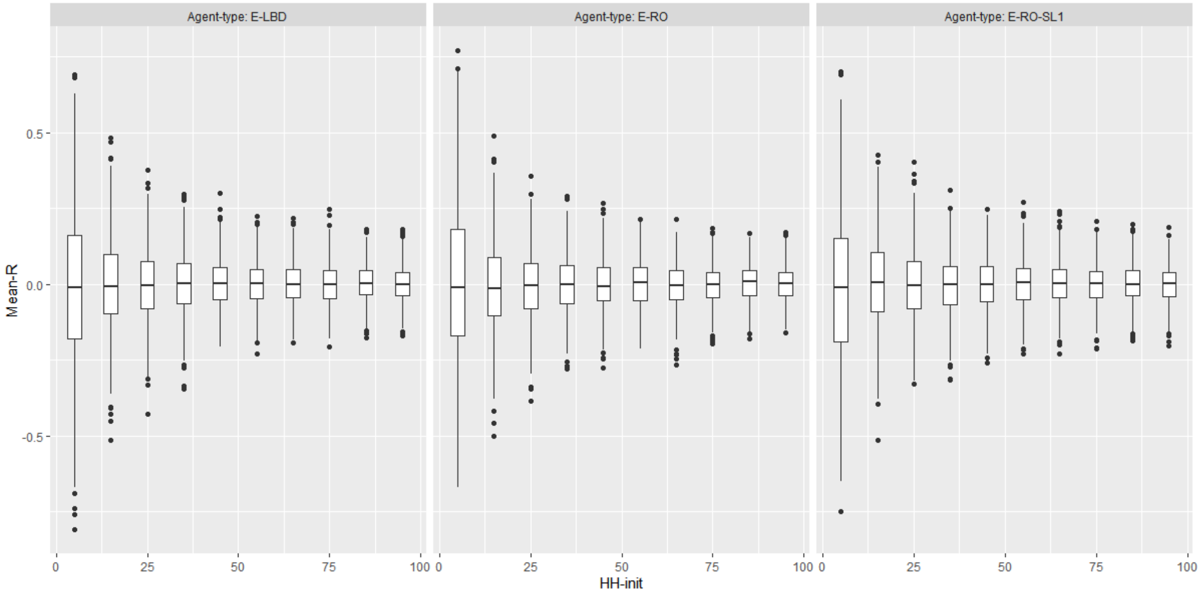

Figure F5.1. Distribution of mean-R values at the beginning of the simulation as a function of the initial total number of households – results for 1000 repetitions for each HH-init value.

F.6. Full sensitivity analysis with 1000 repetitions (90000 runs)

- HH-init: 5 30 50 70 95;
- Livestock-init: 70 75 90 110 160 250;
- E-LBD, E-RO, E-RO-SL1;
- 1000 repetitions.

ECON outcomes as a function of initial number of households (HH-init) and initial number of livestock:

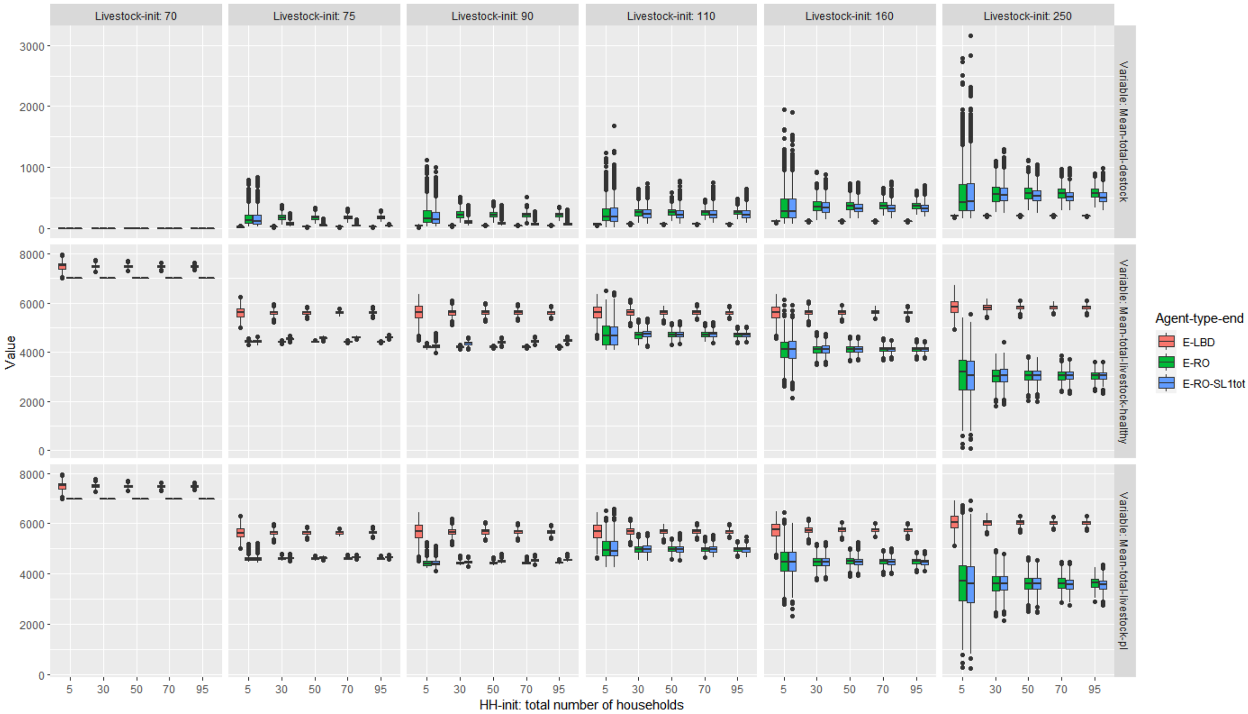

Figure F6.1. Full sensitivity analysis of economic outcomes depending on the initial total number of households and the initial number of livestock.

Zooming in on **mean-total-livestock-healthy**:

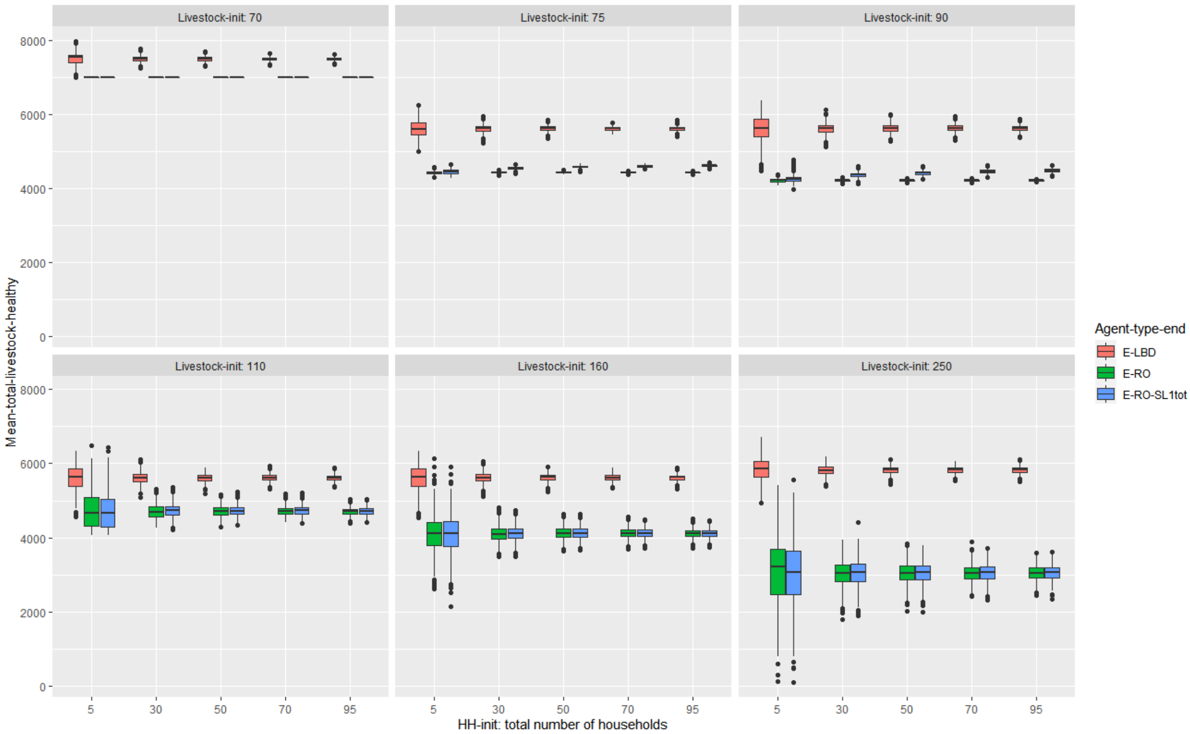

Figure F6.2. Economic results quantified with the mean-total-livestock-healthy – detail from Figure F6.1.

ECOL outcomes as a function of initial number of households (HH-init) and initial number of livestock

Reserve biomass:

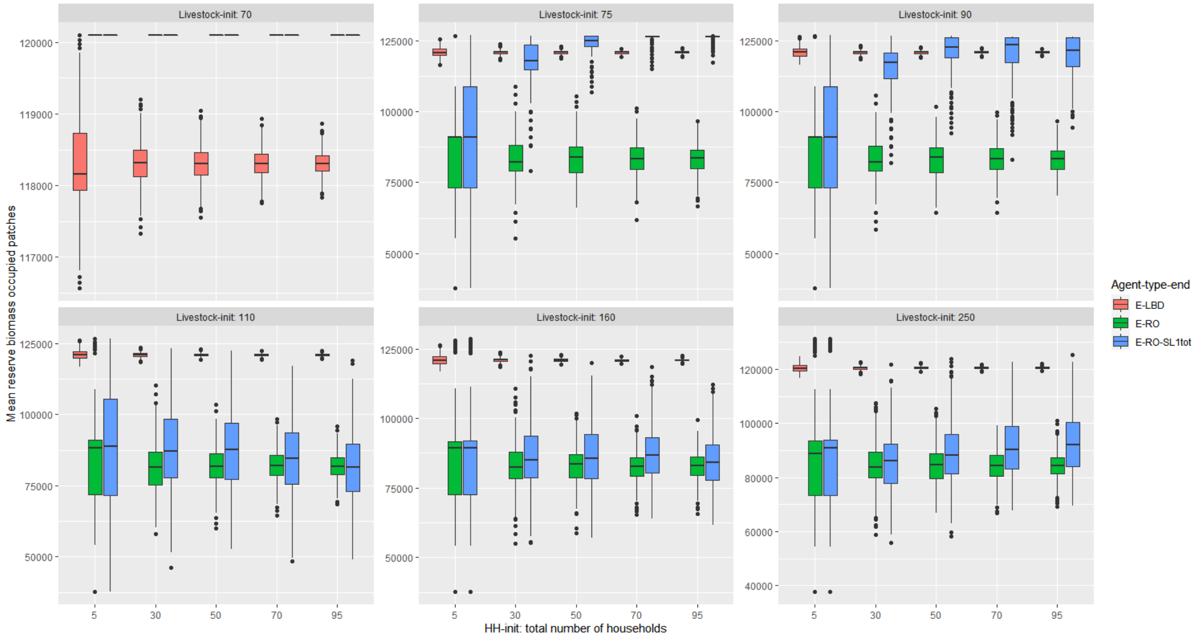

Figure F6.3. Full sensitivity analysis of reserve biomass outcomes as a function of the initial total number of households and the initial number of livestock.

In Figure F6.4, results above are shown with splitting of social learning group between E-SL1 agents who did learn (SL1y) by the end of the simulation and those who did not (SL1n). SL1y agents are agents with the property ***household-SL?*** set to TRUE. This agent property is initialized as FALSE and it is set to TRUE if the agent has engaged in at least one social learning interaction throughout the entire simulation. SL1n are agents whose ***household-SL?*** property was still FALSE at the end of the simulation.

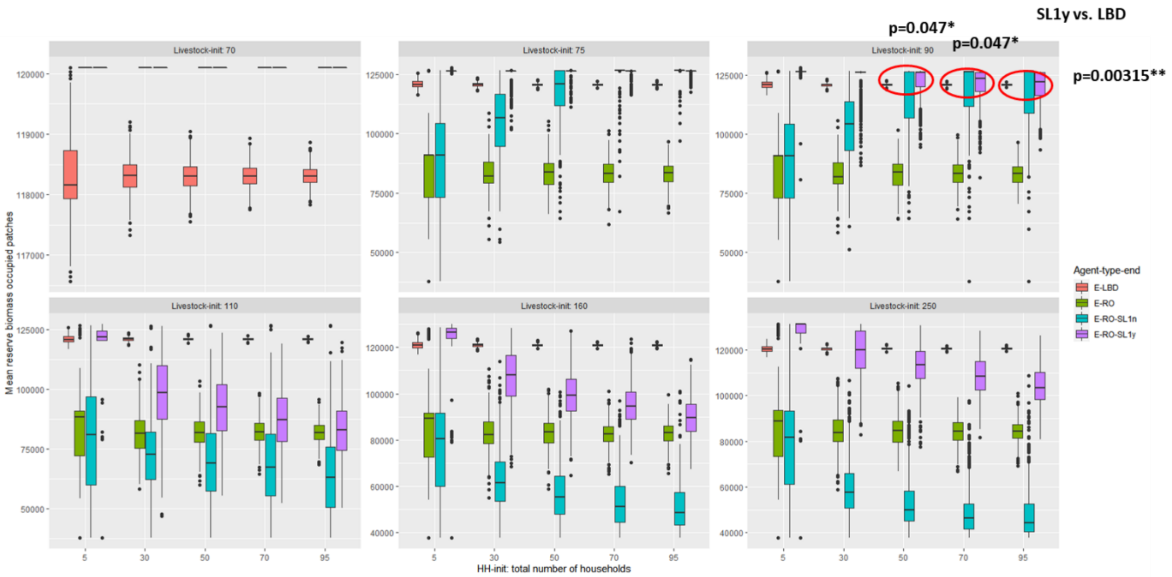

Figure F6.4. Elaboration on Figure F6.3 showing reserve biomass results of social learning 1 agents differentiated by agents who engaged in social learning interactions (SL1y) and those didn't (SL1n).

At moderate initial levels of livestock, biomass is consistently higher for social learning than for learning-by-doing agents. For high initial numbers of livestock (>100), learning-by-doing agents perform better in terms of reserve biomass. However, this is an artefact of the inherent thresholds in the equation of vegetation regeneration (as suggested also in Section F4.), and of the system having an initial carrying capacity of 84 livestock (45000 initial green-biomass + 9000 initial edible reserve biomass -> 51000 / 640 fodder intake = 84.375). It is for this reason that we do not consider results for VERY high initial livestock numbers as relevant. We suggest that future studies test the model in ecological contexts with different carrying capacities.

Green biomass:

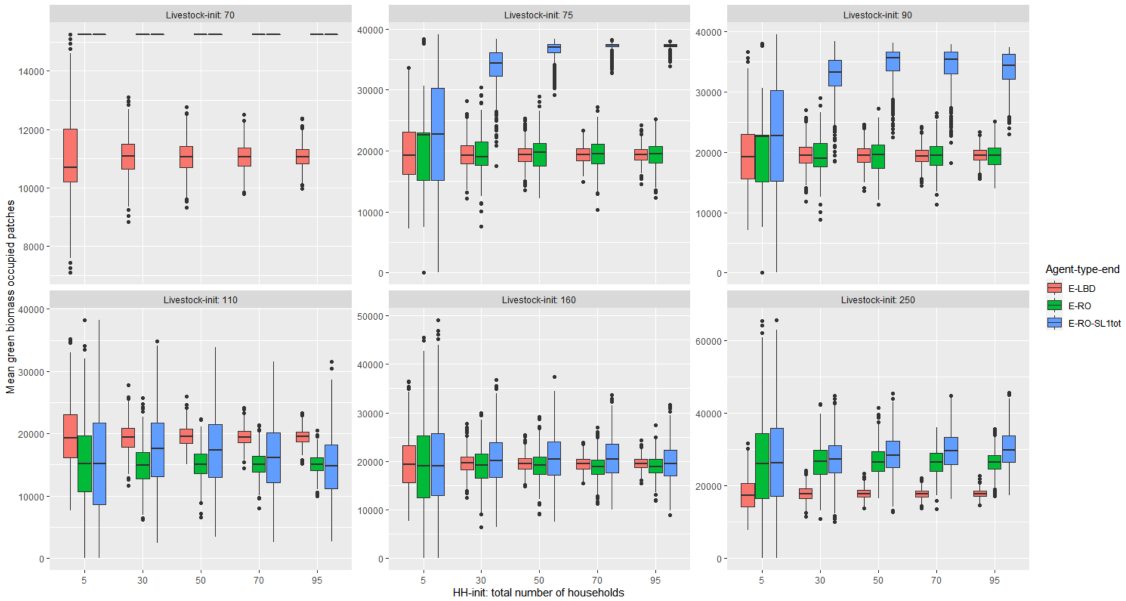

Figure F6.5. Full sensitivity analysis of green biomass outcomes as a function of the initial total number of households and the initial number of livestock.

SOCIAL outcomes - Gini coefficient is calculated for total livestock healthy and 1000 repetitions in each treatment.

The bottom panels of Figure F6.6 show a reversed relation in the agents' type with most equal outcomes compared to the top panels. In the main manuscript we discuss results for livestock-init = 90, where most economic inequality corresponds to learning-by-doing behaviour. Previous analyses have shown that the sensible range of values for the initial number of livestock variable in which to compare learning effects is between 74-100 (see also Figure F1.1).

As discussed above under ECON outcomes, there are inherent constraints in the model related to the carrying capacity which make VERY high livestock-init numbers irrelevant. This is even more clearly underlined by the Gini-index results. Currently, there is no mechanism implemented for baseline agents and social learning agents to account for very high destock numbers. As such, with an initial herd size of 200 and r-value of e.g. -0.95, agents will place 195 livestock on the pasture in the next round. With r-value 0.95, they will place 0. These are more extreme differences (higher Gini-index) than for learning-by-doing agents, which match their decisions more closely to observations of both pasture and underfed livestock.

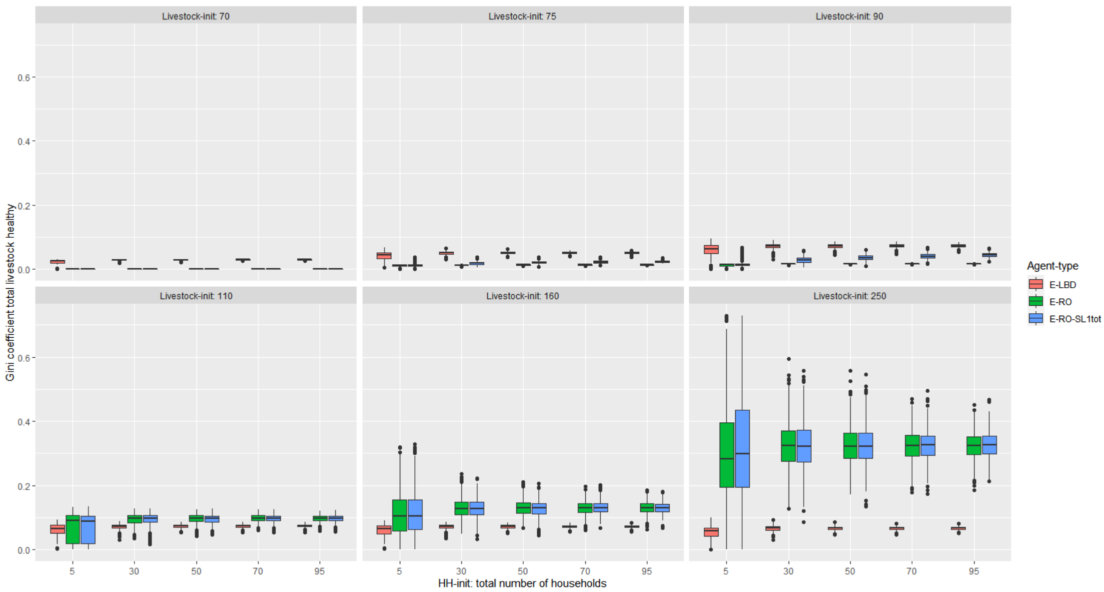

Figure F6.6. Full sensitivity analysis of social outcomes (Gini-coefficient) as a function of the initial total number of households and the initial number of livestock.

F.7. Knowledge radius effects

- 90 livestock-init;
- HH-init: 5, 50, 95;
- 1000 repetitions;
- Knowledge radius: 1, 2, 3.

ECON outcomes:

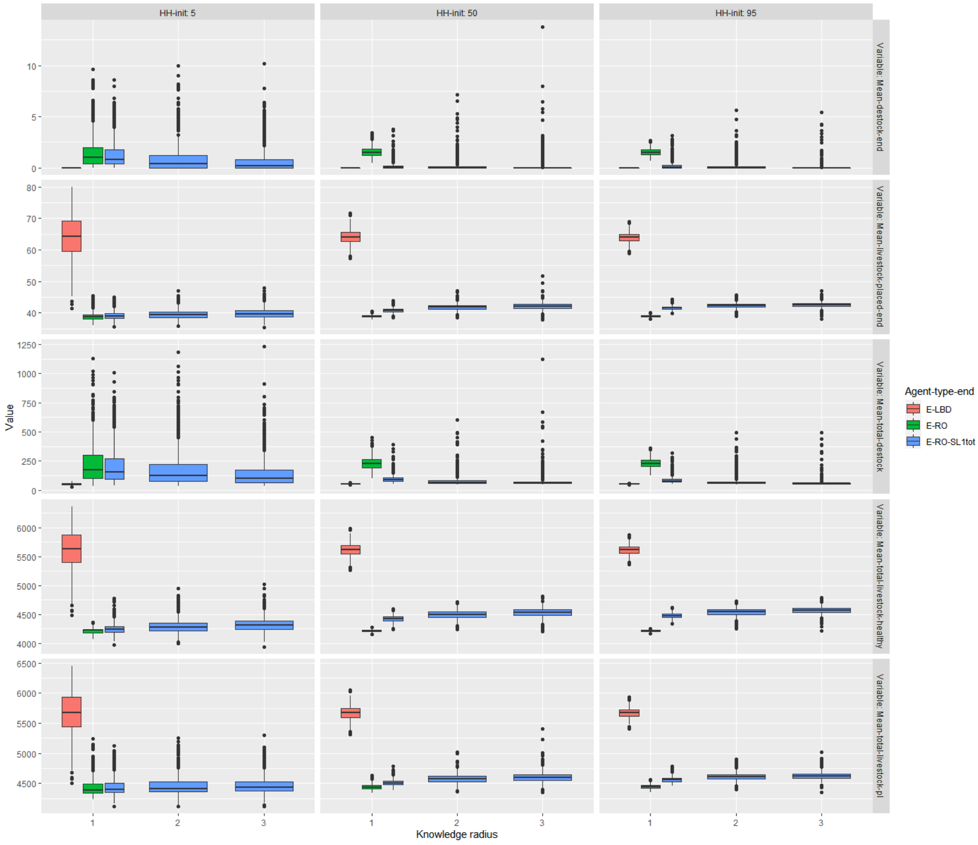

Figure F7.1. Sensitivity analysis of the effect of knowledge radius on economic outcomes.

ECOL outcomes:

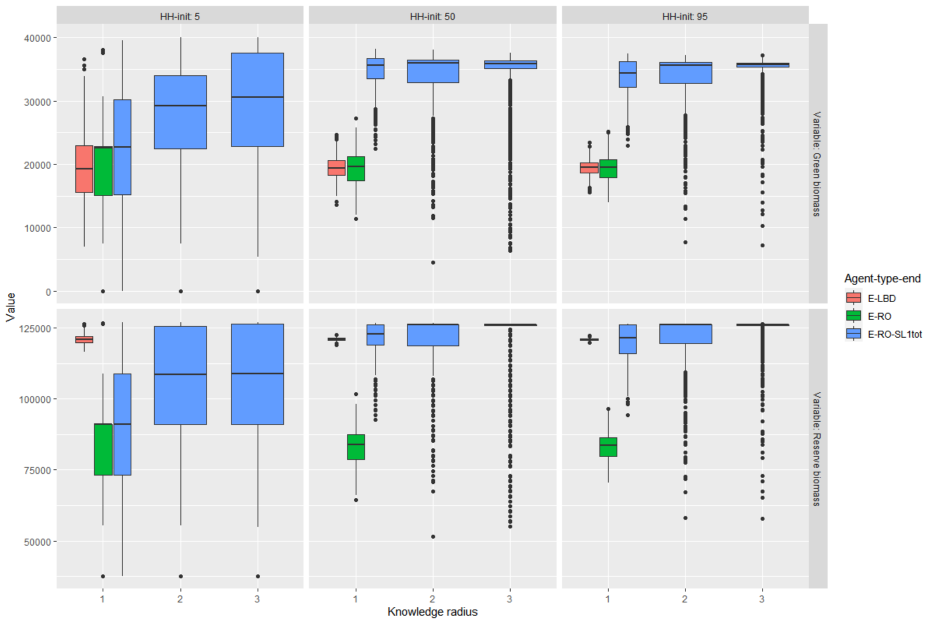

Figure F7.2. Sensitivity analysis of the effect of knowledge radius on ecological outcomes.

SOC outcomes – Gini-coefficient total livestock healthy

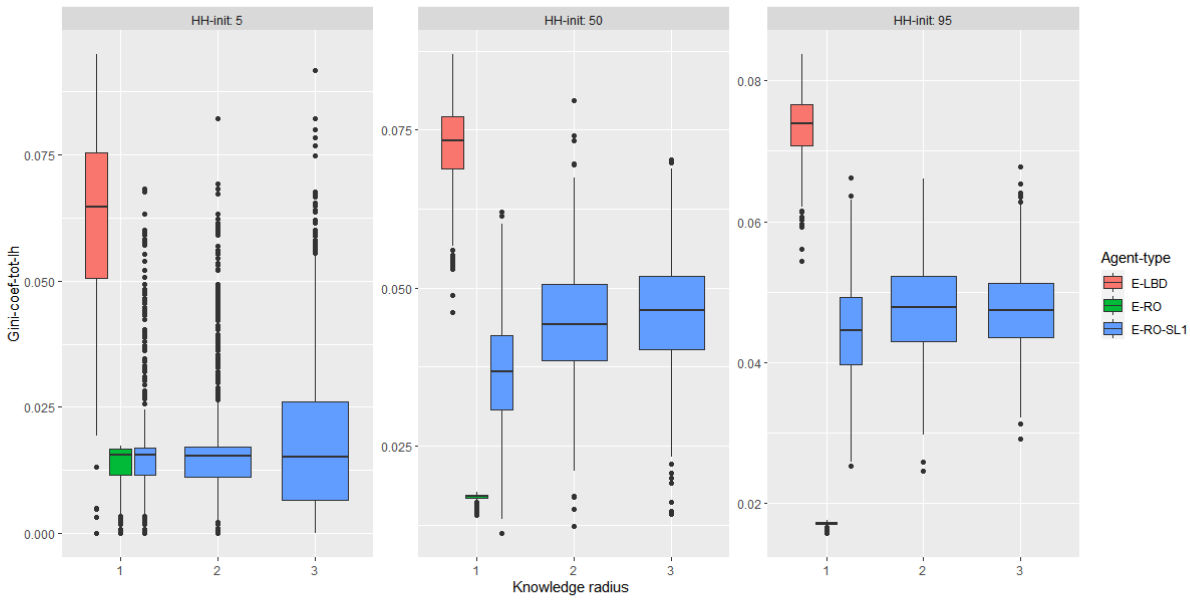

Figure F7.3. Sensitivity analysis of the effect of knowledge radius on social outcomes.

F.8. Outcome variations as a function of Mean-R and HH-init

Below, heatmaps are shown for 50 repetitions, and mean-R (at the beginning of the simulation) is shown binned to values rounded to one decimal. Results are shown for independent simulations with homogeneous populations of each agent type.

ECON results – mean-total-livestock-healthy:

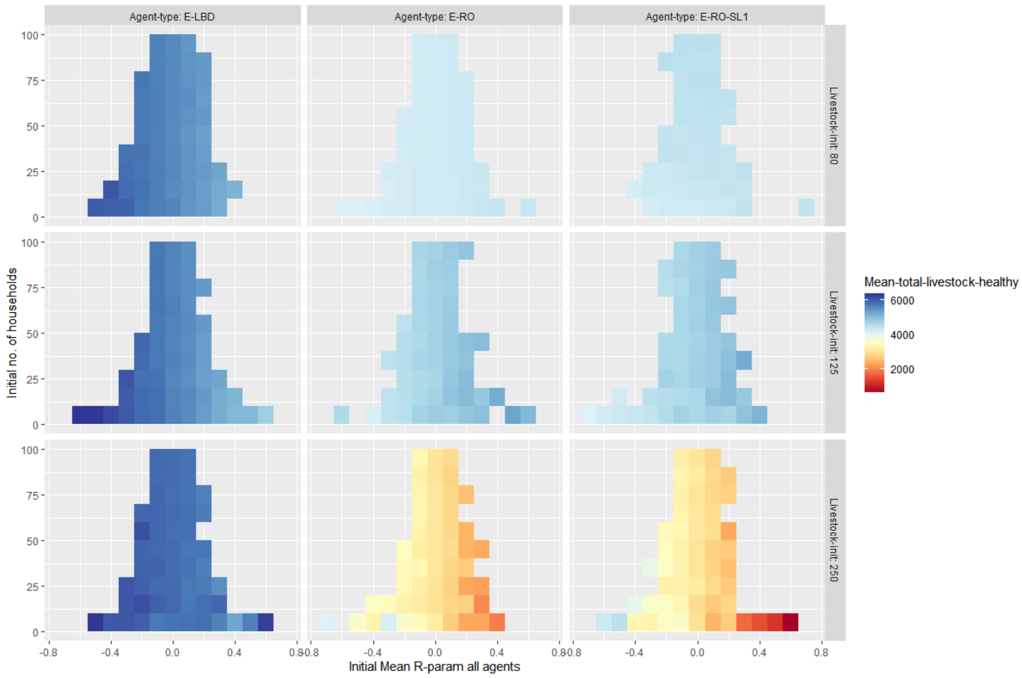

Figure F8.1. Relationship between mean r-parameter values, agent types and economic outcomes (on mean-total-livestock-healthy over all agents at the end of the simulation).

Observations:

- LBD agents perform better for negative R parameters
- E-RO and E-RO-SL1 perform slightly better for higher (and positive) R parameters (e.g. more cautiousness)
- BUT: with increased pressure over the ecosystem (higher Livestock-init), higher R means less speed in adapting to pressure, which leads to overall poorer performance than that of agents with negative R parameter values. This effect is bigger when SL1y is close to zero (because of little agent density, too few HH, e.g. HH-init=5)

ECOL outcomes – reserve biomass:

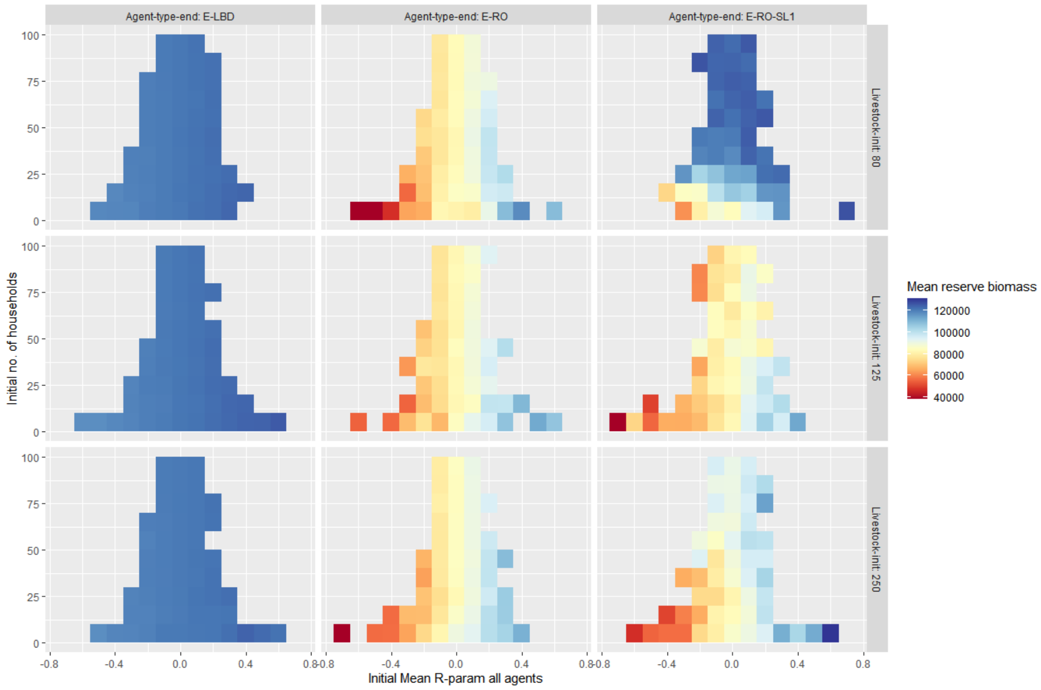

Figure F8.2. Relationship between mean r-parameter values, agent types and reserve biomass.

Green biomass:

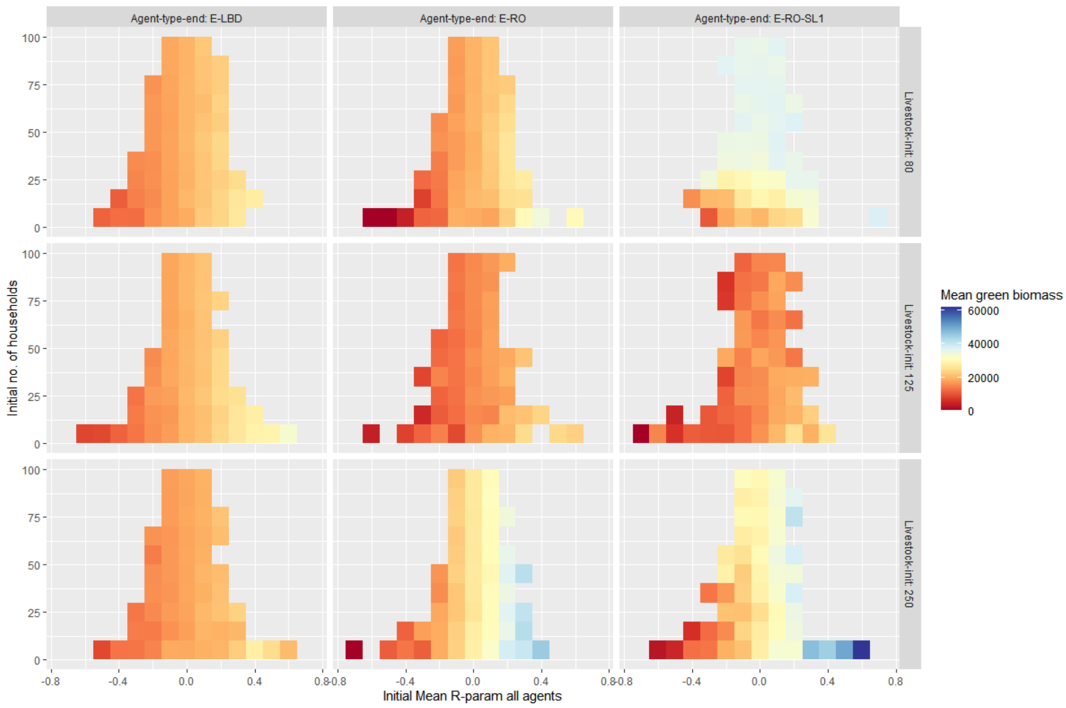

Figure F8.3. Relationship between mean r-parameter values, agent types and green biomass.

SOC outcomes – Gini coefficient for total livestock healthy

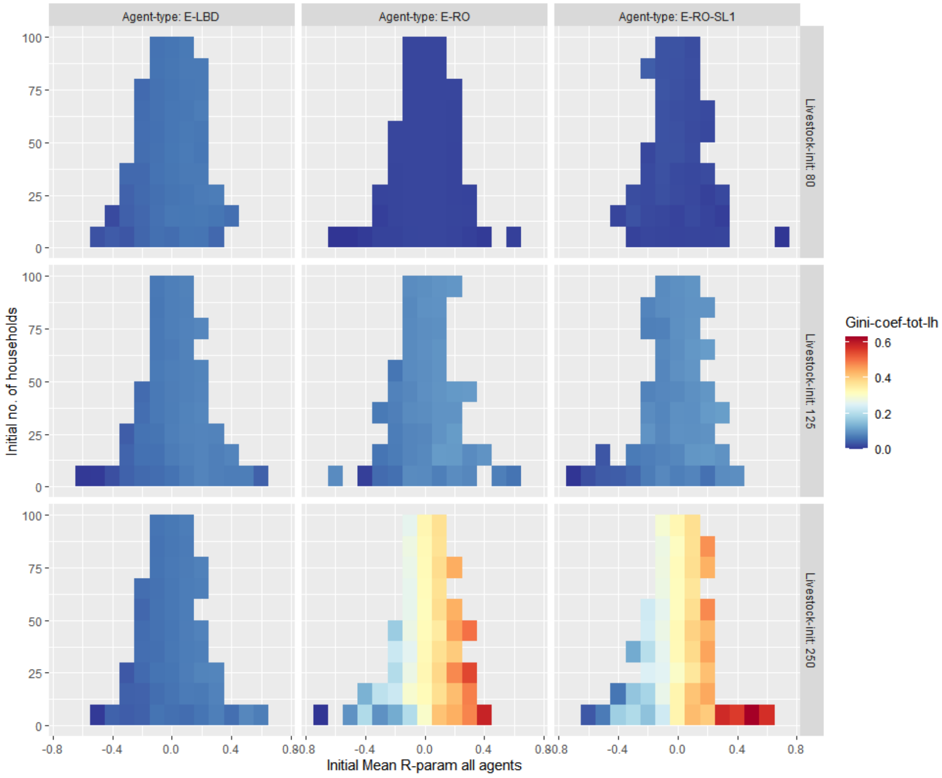

Figure F8.4. Relationship between mean r-parameter values, agent types and social outcomes (mean Gini coefficients over the total number of livestock healthy).

F.9. Changes in Mean-R parameter for Social Learning (E-RO-SL1) agents over time

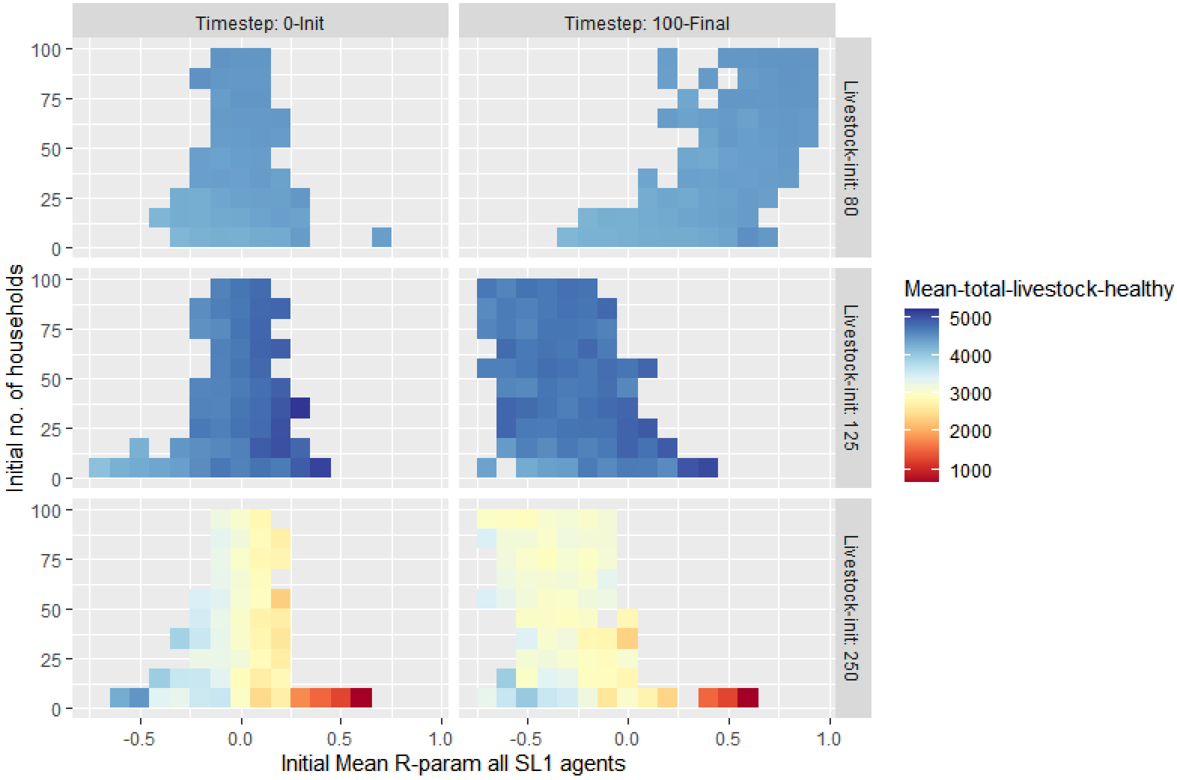

Figure F9.1. Evolution of the distribution of mean-R-parameter in a population of social learning 1 agents (t=0 vs. t=100) and agents' associated economic outcomes depending on their R-parameter values.

At non-extreme initial livestock herd sizes (i.e. <100), social learning agents with positive r-parameter values are more economically successful than those with negative r-parameter values, which results in the more frequent adoption of positive r values.

At very high values of initial livestock herd size, social learning agents will select for negative values of the r-parameter. The carrying capacity of the ecological system is at 84 livestock. The more the initial livestock number is above this value and the r-parameter value is high (for r>0), the more dramatic the reduction in herd size after the first round (due to high destock value). This will lead to poorer economic results than those of other agents, hence over time low values of r-parameter are favoured. Such considerations are also touched upon in the Discussion section of the main manuscript.