DEPARTMENT OF FOOD AND RESOURCE ECONOMICS
UNIVERSITY OF COPENHAGEN

IFRO Working Paper

Modelling land use transition through social learning

Yeqing Duan
Nils Droste
Brian Danley

2026 / 01

IFRO Working Paper 2026 / 01

Modelling land use transition through social learning

Authors: Yeqing Duan, Nils Droste, Brian Danley

JEL-classification: C63, D83, Q24, Q57

Published: April 2026

See the full IFRO Working Paper series here:
https://ifro.ku.dk/english/publications/working_papers/

Department of Food and Resource Economics (IFRO)
University of Copenhagen
Rolighedsvej 23
DK 1958 Frederiksberg DENMARK
https://ifro.ku.dk/english/

Modelling land use transition through social learning

Yeqing Duan[1], Nils Droste[2], and Brian Danley[3]

[1]Environmental Politics Research Group, Department of Political Science, Lund University
[2]Department of Food and Resource Economics, University of Copenhagen
[3]Department of Earth Sciences, Natural Resources and Sustainable Development, Uppsala University

April 27, 2026

Abstract

Land use transition toward multifunctional practices is greatly affected by social learning, yet the temporal interaction between learning mechanisms and network structure remains underexplored. This study examines two social learning channels, information exchange and normative pressure, and how network architecture shapes their effects on transition outcomes. We developed SALT (**S**ocial learning in **A**gent-based **L**and use **T**ransitions), a spatially explicit model that integrates the Consumat framework and reinforcement learning. The model is parameterized using a Swedish forestry context, simulating landowner adaptive decisions under integrated and modular social networks. Results show that the two channels play distinct roles across transition phases. Lack of knowledge limits adoption in early adoption. Individual experience is the main source of knowledge accumulation, and social learning alone cannot close the knowledge gap. As adoption spreads, normative pressure constrains implementation intensity to the prevailing local average, explaining the gap between behavioral and actual landscape changes. Network architecture shapes both channels. Integrated networks widen information exchange and allow alternative-use norms to strengthen over time, while modular networks restrict information circulation and lock in low-implementation local norms. Landscape change organizes along social ties rather than geographic proximity, with architecture determining whether adoption clusters into cohesive blocks or disperses as a diffuse mosaic in the social network. Landowner types contribute differently to behavior change and landscape change across both architectures. These findings suggest that effective transition governance must be tailored

to both phase and social context. Early interventions should prioritize technical assistance, while raising the visible norm of implementation intensity matters more as adoption spreads. In modular communities, consolidating norms within communities before extending outreach is more effective than diffuse seeding. Instruments targeting behavior change need to be paired with those that directly support implementation intensity of alternative practice among less conformity-constrained landowners.

Keywords: Land use transition; Social learning; Social network structure; Agent-based modelling; Multifunctional landscape
JEL Codes: C63, D83, Q24, Q57

1 Introduction

Land systems today face increasing pressure to deliver multiple ecosystem services (ESs). The same land resources are simultaneously expected to produce food, support livelihood, safeguard biodiversity, regulate water cycles, etc (IPCC 2019). These demands are not only growing but often conflicting. For example, timber production competes with other ESs such as carbon sequestration and biodiversity (Blattert et al. 2023). Traditional single-objective land management, optimizing for yield or a single ecosystem service, is increasingly inadequate when societies require diverse ecosystem service portfolios from limited land resources. This multi-objective reality drives the imperative for land use transition [1] toward alternative practices that can deliver diverse ecosystem services, such as multifunctional agriculture and forestry (Hertog, Brogaard, and Krause 2022; Kremen and Merenlender 2018; Terasaki Hart et al. 2023). While policy and institutional support the diffusion of multifunctional land use practices, societal land use transition fundamentally depends on social coordination and learning processes embedded in network structures that shape how information spreads and benefits are perceived (Eastwood et al. 2022; Isaac and Matous 2017; Kremen and Merenlender 2018; Luján Soto et al. 2021).

From both an empirical and a modeling perspective, social network structures are significant conditions for land use transition. At the macro level, the network connectivity affects the potential reach of new information and the distribution of social influence. Weakly connected communities can limit information exchange and diffusion (Phelps, Heidl, and Wadhwa 2012; Loch and Kleinschmit 2025), while integrated communities may reinforce existing norms of single-objective land management by exposing landowners to the majority of conventional practices (Muthukrishna and Schaller 2020; Efferson, Vogt, and Fehr 2020). At the micro level, landowner heterogeneity shapes how these structural conditions translate into individual behavior. Landowners differ in how strongly they weigh private utility versus social conformity (Sotirov, Sallnäs, and Eriksson 2019), which shapes their sensitivity to how knowledge and behavioral expectations propagate through the network. Together, these factors define the social environment through which new land use patterns are built.

Social learning can be both an engine of change and a roadblock in land use transitions, through two mechanisms. First, social ties serve as information channels. They provide technical knowledge as well as the costs and benefits of multifunctional practices (Deutsch and Gerard 1955; Junquera and Grêt-Regamey 2019). This mechanism is less scale-dependent,

[1] In this study, land use transition refers to shifts in land management practices rather than changes in land cover. Management regime changes (e.g., from rotation forestry to continuous cover forestry) alter ecosystem structure and service outputs without necessarily changing the physical extent of a land use type.

as digital platforms allow information to travel through long distances (Skaalsveen, Ingram, and Urquhart 2020). Second, social ties exert normative pressure, where neighbors' attitude and behavior shape the social "appropriateness" of a management shift (Deutsch and Gerard 1955; Junquera and Grêt-Regamey 2019). Unlike information flow, normative influence is more localized as neighbors are typically in close physical proximity (Gu 2024). The relative influence of each mechanism is not uniform across landowners: those primarily oriented toward private returns may be more driven by information exchange, while those who value social expectations may show a greater sensitivity to neighborhood behavior. The interplay between these two mechanisms, filtered through landowner preference heterogeneity, determines whether the social landscape facilitates transition or reinforces the norm of traditional single-objective practice.

Under the right conditions, social learning can gradually transform land use practices across the landscape, eventually leading to the emergence of a new land use pattern. This emergence can be considered a non-linear process. A landscape-level transition may not be self-sustaining unless adoption within the social network reaches a tipping point, at which the adopter community becomes large enough to sustain diffusion pressure across the broader network. Understanding how landscape-level land use patterns emerge from social learning dynamics within communities is therefore essential for identifying context-specific policy leverage points across diverse social contexts (Martin et al. 2022).

Existing studies on the intersection of social networks and land use transition can be grouped into three categories. The first uses surveys and qualitative interviews to reveal how land use choices are affected by social learning and network (e.g., Eastwood et al. 2022; Slijper et al. 2022). The second applies social network analysis and empirical data to explore network properties and influential actors (e.g., Luján Soto et al. 2021; Emami-Skardi, Momenzadeh, and Kerachian 2021; Valujeva et al. 2023). However, these two types are typically static and unable to capture adaptive behavior and dynamic land use shift. Therefore, a lot of studies have employed modeling approaches, such as agent-based modeling (ABM), to simulate land use outcomes under different network conditions (e.g., Manson et al. 2016; Wang et al. 2021; Matous and Bodin 2024). These modeling-based studies often mainly focus on aggregate transition performance metrics (e.g., final adoption rates), rather than on the underlying dynamics of social learning and land use evolution. A knowledge gap remains regarding how land use transition trajectory is shaped by social learning processes, conditional on network structures.

This study develops SALT (**S**ocial learning in **A**gent-based **L**and use **T**ransitions), a spatially explicit ABM for simulating land use transitions shaped by social dynamics. SALT is parameterized with a Swedish forestry context to explore how social learning (informa-

tion exchange and normative pressure) interacts with social network architectures to shape transition dynamics. Specifically: (1) What roles do information exchange and normative pressure play in driving land use transition, and do these roles vary across transition phases and landowner types? (2) How does network architecture shape the operation of these two learning channels and the resulting transition trajectories? (3) What landscape patterns emerge from social learning dynamics, and which landowner types drive these patterns under different network architectures?

2 Methodology

2.1 Conceptual framework

The transition toward multifunctional land use is conceptualized through three interconnected components: social architecture, social learning engine, and spatial landscape (Figure 1).

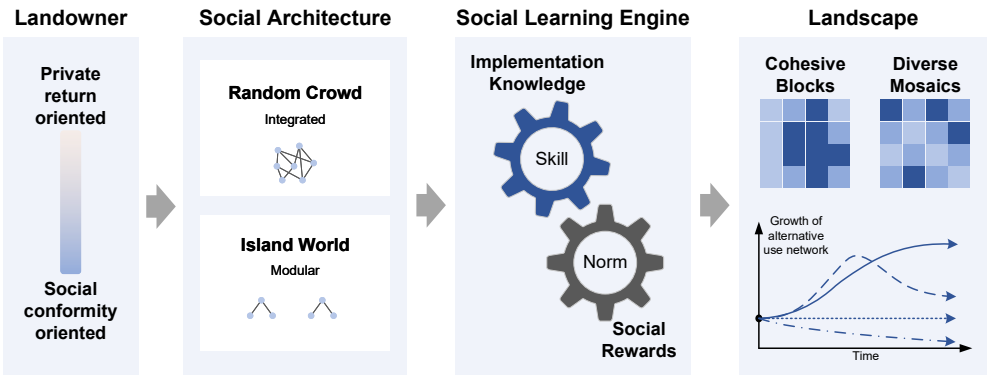

Figure 1: Conceptual framework.

Social architecture defines the structural context within which social learning operates. Network connectivity determines the boundaries of social influence: who exchanges information with whom and whose behavior is visible as a referent. We operationalize this through two contrasting topologies. Random crowd represents an integrated social environment where high connectivity exposes landowners to diverse referents across the network. Island world represents a modular social environment where connectivity is dense within local groups but sparse between them, creating semi-isolated communities with distinct local norms. These two architectures set the structural conditions under which all subsequent learning and norm formation occur.

Social learning engine describes the mechanisms through which landowners update their

knowledge and behavior within a given social architecture. Two channels operate in parallel. The first is information exchange: social ties provide access to technical knowledge about alternative practices, accumulated as implementation knowledge (IK). IK determines whether alternative use is feasible, and grows through both individual experience and social observation. The second is normative pressure: social ties transmit behavioral expectations, captured through social rewards that reflect conformity with the neighborhood majority. Together, these two channels determine not only whether adoption is feasible but also whether it is socially incentivized.

Landowner preference heterogeneity introduces a third conditioning dimension within the social learning engine. Landowners differ in the relative weight they assign to private returns versus social conformity, which determines how strongly each channel drives their decisions. Where private return orientation dominates, information exchange can contribute to private returns of alternative use, largely independent of neighbor behavior. Where social conformity orientation dominates, normative pressure can either enhance behavioral lock-in or accelerate behavioral transition, depending on majority practice in the surrounding community. This heterogeneity does not operate in isolation — it interacts with social architecture. Modular networks that concentrate normative pressure within local groups may activate conformity-oriented landowners, while integrated networks with more global information flows may more effectively reach those constrained by limited implementation knowledge.

The resulting landscape pattern emerges as a consequence of social dynamics between actors within social networks. Geographically, the spatial distribution of alternative use adoption may consolidate into cohesive blocks or remain dispersed as fragmented mosaics depending on how social learning propagates through the network. How these landscape patterns develop conditionally on network architecture and social learning is the main focus of this study.

Together, these three components define the analytical structure guiding this study. Social architecture shapes the conditions under which learning channels operate, which in turn produces distinct adoption dynamics and landscape outcomes.

2.2 Agent-based land use transition model

2.2.1 Model overview

The **S**ocial learning in **A**gent-based **L**and use **T**ransitions (SALT, Figure 2) simulates the dynamics of network evolution underlying the land use transition and explores how the initial structure of the social network affects the land use transition. The model includes N landowner agents and $W \times L$ spatial agents.

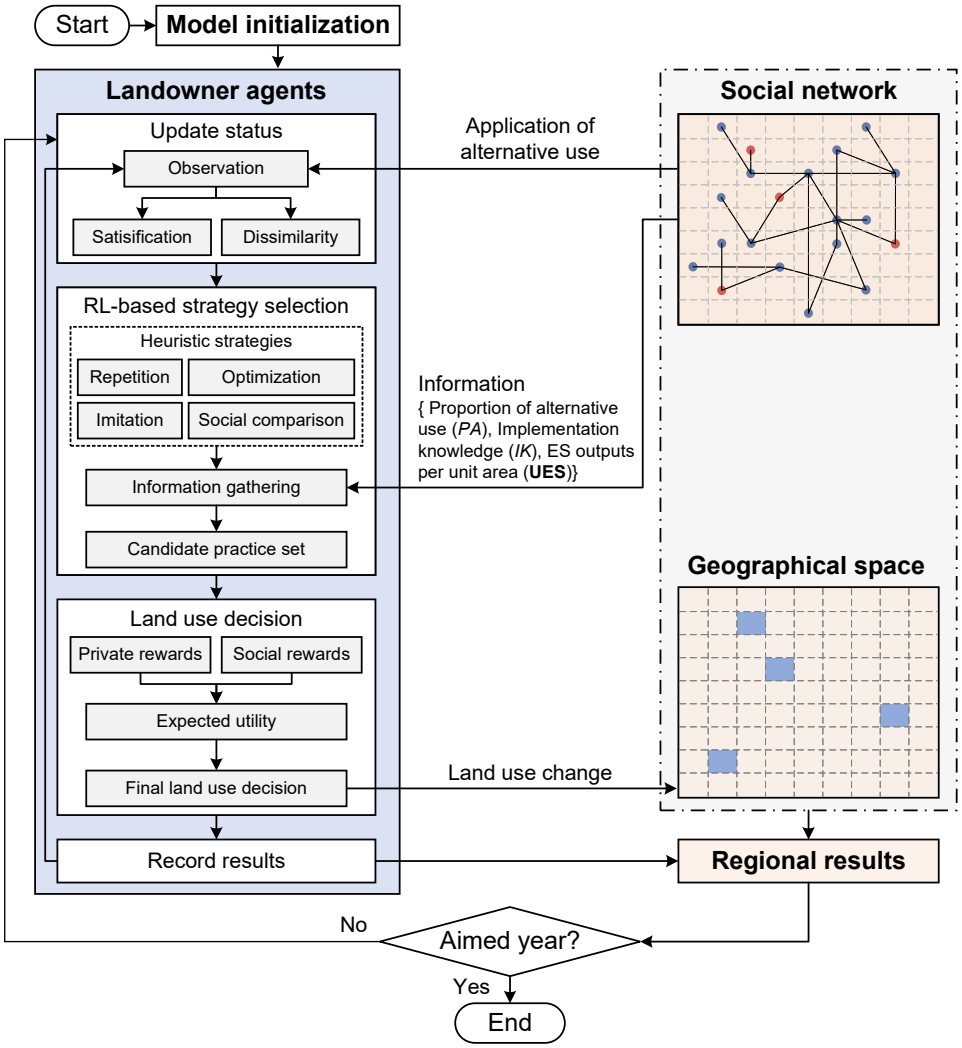

Figure 2: SALT flowchart.

2.2.2 Submodel of ES production

All spatial agents constitute geographical space. All parcels have the same total area but differ in their biophysical productivity (BP_i), reflecting variation in site quality. Each spatial agent represents a land parcel and generates four kinds of ecosystem services (ESs): provision service, carbon sequestration, biodiversity, and recreation.

Two land use regimes are considered: traditional land uses and alternative land uses. Traditional land use aims to maximize economic returns through the production of provision services. The alternative land use aims to diversify ESs and maximize overall ecological value. Any parcel can implement one regime exclusively or both simultaneously (e.g., 80% of the parcel is under traditional use and 20% is under alternative use), consistent with em-

pirical evidence of mixed forestry practices (Stryamets et al. 2026). Each possible allocation combination is referred to as a land use practice. Therefore, land use practices are categorized into pure traditional use and incorporate some alternative use. For simplicity, we adopt the concept of an equivalent annual value, which represents the ES outcomes of each regime as annualized ES outputs. These annualized values serve as performance indicators to compare different land use practices within the model.

The primary goal of this model is to explore the diffusion of alternative use rather than to simulate detailed biophysical ES dynamics. Consequently, the two land use regimes are represented using annualized ES outputs per unit area: provision ($UES_{1,i}^{\text{trad/alt}}$), carbon stock ($UES_{2,i}^{\text{trad/alt}}$), biodiversity ($UES_{3,i}^{\text{trad/alt}}$), and recreation ($UES_{4,i}^{\text{trad/alt}}$) to reduce complexity. The superscript "trad/alt" indicates that each ES output is defined separately for the traditional ("trad") and alternative ("alt") regimes. These annualized indicators summarize the long-term average ES performance of each regime, providing a simplified yet realistic representation of ES supply. The overall unit output of each ES ($UES_{e,i,t}$) on parcel i at time t is calculated as the area-weighted sum of the outputs from the two land use regimes.

2.2.3 Submodel of landowner agents

Each landowner agent owns a single parcel, which can simultaneously allocate different fractions of its area to traditional and alternative uses.

All landowner agents constitute the whole social network. The simulated social network consists of two levels. The first layer is the personal network of each landowner agent (local level), comprising the agent and its direct social connections. The second layer is the entire simulated social network (global level). Landowner agents adapt their behavior to pursue utility maximization (simulation duration and landowner agents' lifespan?). The total utility consists of private rewards derived from annualized ES outputs and social rewards obtained through the agent's network. To capture different preferences on ESs and social rewards, the model considers three types of landowner agents: economic-oriented, multiobjective, and amenity-oriented.

By integrating the Consumat approach (Janssen, Vlek, and Jager 1999; Jager et al. 2000), its variation (Huber et al. 2022), and reinforcement learning, landowner agents follow an "observation-strategy-decision" decision-making process in our model. The decision process has three steps.

2.2.3.1 Update status

At the beginning of each simulation round, each landowner agent observes its historical results as well as the application of alternative uses within its personal network. Following Huber et al. (2022), landowner agents calculate a satisfaction index $SI_{i,t}$ and two dissimilarity indexes (performance $DI^{\mathrm{p}}_{i,t}$ and behavior $DI^{\mathrm{b}}_{i,t}$). These metrics collectively determine the agent's satisfaction status and social orientation.

Adapted from Duinen et al. (2016) and Kreft et al. (2024), $SI_{i,t}$ is the ratio between the previous private utility (i.e., weighted sum of the annualized four ES outputs under current land use practice, see in section 2.2.3.3) $U^{\mathrm{private}}_{i,t-1}$ and the average private utility $\bar{U}^{\mathrm{private}}_{i,t-1}$ over landowner agent's memory length ML_i, as shown in Eq. (1). A higher $SI_{i,t}$ implies the landowner agent is more satisfied with its current land use practice.

$$SI_{i,t} = \frac{U^{\mathrm{private}}_{i,t-1}}{\bar{U}^{\mathrm{private}}_{i,t-1}} = \frac{U^{\mathrm{private}}_{i,t-1}}{\frac{1}{ML_i}\sum_{m=1}^{ML_i} U^{\mathrm{private}}_{i,t-m}} \tag{1}$$

To assess compliance with social norms, landowner agents utilize two dissimilarity indices based on Huber et al. (2022). The behavioral deviation index ($DI^{\mathrm{b}}_{i,t}$) quantifies the difference between the agent's alternative use ratio and the personal network average ($\overline{PA}^{\mathrm{Nei}}_{i,t}$), as defined in Eq. (2). Performance dissimilarity index ($DI^{\mathrm{p}}_{i,t}$) is designed as an asymmetric metric that assesses underperformance relative to the average person network performance ($\overline{UES}_{e,i,t}$), standardized by biophysical productivity (Eqs. 4–5). By focusing on negative deviations, $DI^{\mathrm{p}}_{i,t}$ reflects a social comparison logic where only inferior ES outcomes trigger behavior changes. Higher values for either index indicate a greater divergence from the prevailing social norm.

$$\overline{PA}^{\mathrm{Nei}}_{i,t} = \frac{1}{N^{\mathrm{Nei}}_i}\sum_{n=1}^{N^{\mathrm{Nei}}_i} PA^n_{i,t} \tag{2}$$

$$DI^{\mathrm{behave}}_{i,t} = \frac{\left|PA_{i,t} - \overline{PA}^{\mathrm{Nei}}_{j,t}\right|}{\max(PA_{i,t}, \overline{PA}^{\mathrm{Nei}}_{j,t})} \tag{3}$$

$$\overline{UES}_{e,i,t} = \frac{1}{N^{\mathrm{Nei}}_i}\sum_{n=1}^{N^{\mathrm{Nei}}_i} UES_{e,n,t} \tag{4}$$

$$DI^{\mathrm{perform}}_{i,t} = \sum_{e=1}^{4} w_{e,i}\frac{\max\left(\overline{UES}_{e,i,t}/\overline{BP} - UES_{e,i,t}/BP_i, 0\right)}{\max\left(UES_{e,i,t}/BP_i, \overline{UES}_{e,i,t}/\overline{BP}\right)} \tag{5}$$

where $PA^n_{i,t}$ refers to the proportion of alternative use applied by peer n of agent i at time

t; N_i^{Nei} is the number of peers connected to landowner agent i; $PA_{i,t}$ denotes the proportion of alternative use applied by agent i at time t; $UES_{e,n,t}$ is the e-th annualized ES outputs per unit area of peer n at time t; $w_{e,i}$ is the landowner agent's preference of e-th ES.

2.2.3.2 RL-based strategy selection

The landowner agent selects a heuristic decision strategy (repetition, optimization, imitation, and social comparison) based on its satisfaction and dissimilarity status. The action rules of the four heuristic strategies are as follows. (i) Repetition: the landowner agent simply continues its previous land use practice. (ii) Optimization: the landowner agent seeks to maximize its utility based on available information and individual experimentation. (iii) Imitation: the landowner agent copies a practice applied within its personal network to improve its utility. (iv) Social comparison: the landowner agent collects information both within and outside its personal network and chooses the practice that can maximize its overall utility. These strategies determine the subsequent choice of a specific land-use practice.

Reinforcement Learning (RL) is applied in this study to enable agents to dynamically choose the optimal strategy given their satisfaction and dissimilarity status. The traditional Consumat approach and its variations employ fixed thresholds to externally predefine the strategy selection rule. For example, a satisfied and individually-oriented agent adopts the repetition strategy in the Consumat model. However, this method relies only on the real-time status, resulting in short-sighted behavior. In contrast, the RL framework aims to optimize the long-term accumulated payoff by proactively learning which strategy yields the best outcome under given status (Hung and Yang 2021; Sanderson, McAllister, and Helgeson 2025; Zhang, Valencia, and Chang 2023). That is, agents can learn the optimal learning rules (i.e., when to switch behavioral strategies) based on RL. Therefore, RL offers a more realistic simulation of landowner behavior, as land use and management decisions inherently require long-term deliberation rather than instantaneous, short-term reactions.

To enable landowner agents to learn optimal strategies adaptively based on their satisfaction and dissimilarity status, this study introduces the Deep Q-Network (DQN) algorithm. DQN is a value-based RL algorithm that integrates Q-learning with a Neural network (NN) to handle continuous state spaces.

(1) Strategy selection process based on DQN

(i) RL core elements and function approximation

The RL learning environment is defined by the State ($s_{i,t}$), Actions ($a_{i,t}$), and Reward ($U_{i,t}$) for agent i. **State** ($s_{i,t}$): Each landowner agent i at time t is characterized by a continuous state vector $s_{i,t} = (SI_{i,t}, DI_{i,t}^{\text{p}}, DI_{i,t}^{\text{b}})$. **Actions** ($a_{i,t}$): The landowner agent selects

one of four discrete Consumat strategies ($a_{i,t}$) which form the action space $a_{i,t} \in \mathcal{A} =$ {Repetition, Optimization, Imitation, Social comparison}. **Reward** ($U_{i,t}$): The reward is the instant utility ($U_{i,t}$) obtained after executing a strategy, serving as the feedback signal for learning. Both private and social rewards are considered for utility calculation (details in section 2.2.3.3).

In classical Q-learning, the action-value function $Q(s_{i,t}, a_{i,t})$ is used to estimate the expected long-term cumulative reward of executing strategy $a_{i,t}$ in state $s_{i,t}$, as shown in Eq. (6):

$$Q(s_{i,t}, a_{i,t}) = \mathbb{E}\Big[\sum_{m=0}^{\infty} \gamma^m U_{i,t+m} \Big| s_{i,t}, a_{i,t}\Big], \tag{6}$$

where $U_{i,t}$ is the utility (reward) obtained at time t and $\gamma \in (0, 1)$ is the discount factor weighting future relative to immediate rewards.

In the Q-learning algorithm, every possible $< s_{i,t}, a_{i,t} >$, is stored in the Q-table. When state space and action space are large, the Q-table grows quickly and becomes computationally infeasible for the continuous state space in our model. Instead, DQN uses a NN, parameterized by θ, to approximate this action-value function as:

$$Q(s_{i,t}, a_{i,t}; \theta) \approx \mathbb{E}\Big[\sum_{m=0}^{\infty} \gamma^m U_{i,t+m} \mid s_{i,t}, a_{i,t}\Big]. \tag{7}$$

(ii) Strategy selection

Landowner agents select strategies using an ϵ-greedy policy to balance exploitation of learned knowledge and exploration of new options:

$$a_{i,t} = \begin{cases} \arg\max_a Q(s_{i,t}, a; \theta), & \text{with probability } 1-\epsilon \text{ (Exploitation);} \\ \text{random choice,} & \text{with probability } \epsilon \text{ (Exploration).} \end{cases} \tag{8}$$

A higher ϵ encourages exploration of new strategies, while a lower ϵ promotes the exploitation of previously learned, high-value behavior. Following standard DQN practice, ϵ decays exponentially over training, shifting agents from broad exploration toward exploitation of learned Q-value estimates.

(iii) Experience replay and network update

After executing the chosen strategy $a_{i,t}$ and the final land use practice, the transition tuple $(s_{i,t}, a_{i,t}, U_{i,t+1}, s_{i,t+1})$ is stored in an experience replay buffer ($\mathcal{D}$). Agents of the same type share Q-network weights, on the assumption that similar value orientations imply comparable decision logic. Experiences are kept separate as each agent's state reflects its unique network position, neighborhood composition, and biophysical context. The shared network

is updated by each agent drawing mini-batches from its own buffer. The DQN then updates the Q-network (θ) by minimizing the temporal-difference (TD) loss over these sampled mini-batches. The loss function is defined as:

$$\mathcal{L} = \mathbb{E}_{\mathcal{D}}\Big[\Big(U_{i,t} + \gamma \max_{a'} Q(s_{i,t+1}, a'; \theta^{-}) - (s_{i,t}, a_{i,t}; \theta)\Big)^2\Big], \tag{9}$$

where a' represents possible actions at the next state $s_{i,t+1}$, and θ^{-} are the parameters of a target network, which is periodically synchronized with the main Q-network (θ) to stabilize learning.

Through iterative interactions with the evolving environment, agents learn to associate state patterns ($SI_{i,t}, DI_{i,t}$) with optimal strategy choices, thereby capturing the adaptive process of land use transition.

(2) Information gathering and candidate practice set

For maximization, imitation, and social comparison strategies, landowner agents need to collect information and form a set of candidate practices ($\mathcal{B}_{i,t}$) according to the corresponding action rule.

Optimization. Under the optimization strategy, landowner agents seek to improve their utility via bounded innovation. Innovation is operationalized as a localized search around an "anchor" value that reflects the agent's experience with alternative use. The anchor is defined as (i) the current implementation ratio if the alternative use is presently applied, (ii) the most recent implementation ratio if the alternative use was applied in the past, or (iii) zero if the agent has no prior experience. Given a search width sw, the landowner agent draws two values from a normal distribution $N(0, sw)$, which are added to the anchor to generate perturbed alternatives. Specifically, when the anchor is zero, the landowner agent needs to draw two positive values from $N(0, sw)$. Together with the anchor itself, these form a set of three candidate practices.

Imitation. Under the imitation strategy, the landowner agent seeks to improve its utility by imitating the practices of its peers. To do so, the landowner agent contacts its direct connections and collects information on their practices. For connection n, the landowner agent obtains a bundle of information consisting of: the proportion of alternative use (PA_t^n), implementation knowledge (IK_t^n), ES outputs per unit area of alternative use ($\mathbf{UES}_{n,t}^{\text{alt}} = \{UES_{1,n,t}^{\text{alt}}, UES_{2,n,t}^{\text{alt}}, UES_{3,n,t}^{\text{alt}}, UES_{4,n,t}^{\text{alt}}\}$).

Social comparison. Under the social comparison strategy, landowner agents seek to enhance their status by contacting both their peers and individuals outside their personal networks (without establishing new connections). Information gathering from contacted agents follows the same procedure as in the imitation strategy. However, the likelihood

of contacting strangers is not uniform. Previous research has pointed out that homophily plays a significant role in our communication, whereby individuals are more likely to interact with others who share similar characteristics. Building on these studies, three dimensions of homophily are considered in the model: geographical proximity, common connections (Dong et al. 2022), and property similarity (measured by the Euclidean distance, (Wolf et al. 2015)).

Geographical distance is important in shaping local interactions, which is more common than non-local interactions even in the presence of modern communication technologies (Spencer 2012). Geographical distance is also found to be critical to the diffusion of land use practices (Evans et al. 2024). Therefore, geographical proximity ($GP_{i,j}$) affects the likelihood of a focal agent contacting strangers and is calculated as:

$$GP_{i,j} = \frac{1}{1 + d_{i,j}}, \tag{10}$$

where $d_{i,j}$ ($d_{i,j} > 0$) is the geographical distance between agent i and j.

Common connections are important bridges that enable two otherwise unacquainted landowners to become aware of each other. This is measured by the structural similarity of their personal networks $NS_{i,j}$ (Dong et al. 2022), as shown in Eq. (11):

$$NS_{i,j} = \frac{NM_{i,j}}{\max(N_i^{\mathrm{Nei}}, N_j^{\mathrm{Nei}})}, \tag{11}$$

where $NM_{i,j}$ is the number of common connections.

Property similarity refers to the similarity of parcels' biophysical productivity (BP), which strongly influences the ES outputs of given land use practices and thus provides a significant reference for land management. Property similarity $PS_{i,j}$ is quantified using the Euclidean distance as shown in Eq. (12).

$$PS_{i,j} = 1 - \left| \frac{BP_i - BP_j}{\max BP} \right|. \tag{12}$$

The overall homophily, $H_{i,j}$, is calculated using Eq. (13). Since the range of $GP_{i,j}$ is (0, 1), the value of $H_{i,j}$ is strictly bounded within (0, 1).

$$H_{i,j} = \frac{1}{3}(GP_{i,j} + NS_{i,j} + PS_{i,j}). \tag{13}$$

2.2.3.3 Land use decisions

After selecting a Consumat strategy ($a_{i,t}$) via DQN, the specific land use decision ($b_{i,t} \in \mathcal{B}_{i,t}$) is determined according to the strategy's action rule. Specifically, the landowner agent

repeats its previous land use practice if the repetition strategy is selected; otherwise, it chooses a practice from the candidate practice set ($\mathcal{B}_{i,t}$) following the utility maximization rule.

The expected utility ($U_{i,t}$) of each candidate practice is defined as the weighted sum of expected private ($U_{i,t}^{\text{private}}$) and social rewards ($U_{i,t}^{\text{social}}$) and is obtained via Eq. (14):

$$U_{i,t} = (1 - \varphi_{i,t})U_{i,t}^{\text{private}} + \varphi_{i,t}U_{i,t}^{\text{social}}, \tag{14}$$

where $\varphi_{i,t} \in [0, 1]$ denotes the weight of social rewards and is determined according to the agent type. The value of $\varphi_{i,t}$ is determined externally according to the agent type.

Social rewards ($U_{i,t}^{\text{social}}$) measure the assimilative forces originating from the agent's personal network. Specifically, $U_{i,t}^{\text{social}}$ aims to promote behavioral conformity with the majority, measured by the absolute difference between the agent's alternative use proportion and the average within its personal network.

$$U_{i,t}^{\text{social}} = \kappa \left(1 - \left|PA_{i,t} - \overline{PA}_{i,t}^{\text{nei}}\right|\right), \tag{15}$$

where κ reflects the maximum social rewards that the agent can obtain.

Private rewards ($U_{i,t}^{\text{private}}$) are calculated based on the four expected annualized ES outputs ($\mathbb{E}[ES_{e,i,t}]$) and the agent's ES preferences ($w_{e,i}$).

$$U_{i,t}^{\text{private}} = \sum_{e=1}^{4} w_{e,i} \ln \mathbb{E}[ES_{e,i,t}]. \tag{16}$$

The expected total annualized output of each ES ($\mathbb{E}[ES_{e,i,t}]$) is determined by outputs from alternative and traditional uses:

$$\mathbb{E}[ES_{e,i,t}] = area_i PA_{i,t}\mathbb{E}[UES_{e,i,t}^{\text{alt}}] + area_i(1 - PA_{i,t})UES_{e,i,t}^{\text{trad}}, \tag{17}$$

where $area_i$ is the managed land area of agent i; $PA_{i,t}$ is the proportion of land allocated to the alternative use; $\mathbb{E}[UES_{e,i,t}^{\text{alt}}]$ is the expected annualized unit output of ES e of the alternative use; $UES_{e,i,t}^{\text{trad}}$ denotes the annualized unit output of ES e of traditional use and is known by the landowner agent.

Both the expected ($\mathbb{E}[UES_{e,i,t}^{\text{alt}}]$) and actual ($UES_{e,i,t}^{\text{alt}}$) annualized ES outputs of the alternative use are depend on the landowner agent's implementation knowledge ($IK_{i,t} \in [0, 1]$), which represents the ability to realize the full potential of the alternative use. $IK_{i,t}$ gets updated under two conditions: the landowner agent (1) is engaging in alternative use and (2) collects information from other agents (under the imitation and social comparison strategy).

The model assumes that landowner agents can learn by doing. After each land use

decision, if the landowner agent is implementing alternative use, its $IK_{i,t}$ updates based on an update rate l, which is randomly drawn from zero to the learning rate l^*, as shown in Eq. (18):

$$IK_{i,t} = \min[1, (1-l)IK_{i,t} + lIK^{\max}] = \min[1, (1-l)IK_{i,t} + l]. \tag{18}$$

Under the imitation and social comparison strategies, landowner agents update their $IK_{i,t}$ based on collected information and the update rate l (randomly drawn from zero to the learning rate l^*), as shown in Eq. (19):

$$IK_{i,t} = \begin{cases} \min[1, (1-l)IK_{i,t-1} + l(IK^*_{j,t-1} - IK_{i,t-1})], & IK^*_{j,t-1} > IK_{i,t-1}; \\ IK_{i,t-1}, & IK^*_{j,t-1} \leq IK_{i,t-1}, \end{cases} \tag{19}$$

where $IK^*_{j,t-1}$ is the highest knowledge level within the practice set ($\mathcal{B}_{i,t}$).

Higher $IK_{j,t}$ reduces estimation uncertainty and enhances implementation efficiency. The model assumes a linear relationship between $IK_{j,t}$ and ES outputs. Accordingly, the expected and actual annualized unit ES outputs of alternative use can be obtained with Eq. (20) and Eq. (21), respectively:

$$\mathbb{E}[UES^{\text{alt}}_{e,i,t}] = \overline{UES}^{\text{alt}}_{e,i,t}[1 - B_0(1 - IK_{i,t})\varepsilon^0_{i,t}], \tag{20}$$

$$UES^{\text{alt}}_{e,i,t} = min(UES^{\text{alt},*}_{e,i,t}[1 - B_0(1 - IK_{i,t})]\varepsilon^1_{i,t}, UES^{\text{alt},*}_{e,i,t}), \tag{21}$$

where $\overline{UES}^{\text{alt}}_{e,i,t}$ denotes the average observed e-th annualized unit ES output of the alternative use; $UES^{\text{alt},*}_{e,i,t}$ is the theoretical maximum e-th annualized unit ES output of the alternative use; B_0 is the maximum estimation/implementation bias and determines the worst ES outputs; $\varepsilon^0_{i,t}, \varepsilon^1_{i,t} \in [0.9, 1.1]$ represent the random risk.

$\overline{UES}^{\text{alt}}_{e,i,t}$ is initially set to the global average ES output of the entire model. Thereafter, $\overline{UES}^{\text{alt}}_{e,i,t}$ is updated based on both the agent's own implementation experience and information obtained from other agents. When evaluating candidate practices, the landowner agent forms a new observation of the average unit ES output of the alternative use through computing the average of observed unit ES outputs, including its own and those obtained from the information bundle (if not available, use its own experience as the new observation). $\overline{UES}^{\text{alt}}_{e,i,t}$ is then updated as the moving average of the latest observations over the agent's memory length (including the newly added observation). If the agent has neither implemented alternative use nor acquired information from other agents, $\overline{UES}^{\text{alt}}_{e,i,t}$ will retain its previous value.

The selected land use practice will be implemented and cause land use changes to parcels. After implementation, the realized ES outputs and social rewards are observed and used to calculate the actual utility and update the DQN-based strategy selection.

2.2.4 Simulation outputs and analytical framework

As landowner agents adopt alternative practices, they collectively form an alternative use network that evolves endogenously throughout the simulation. At each step, the network expands or contracts as agents update their land use decisions, with IK and social rewards jointly shaping adoption and reversion outcomes. The resulting landscape pattern emerges from the bottom-up interaction of individual decisions within a given social architecture. Its structure and trajectory reflect how the two social learning channels operate and propagate through the social network over time. The analysis follows the three conceptual components of the framework to examine how social architecture shapes learning dynamics and produces distinct landscape outcomes.

Social architecture effects. Final adoption outcomes and characteristics of the emerged alternative use network are used to assess the effect of network architecture. Specifically, final adoption outcomes include both adoption rate and alternative use area (% of total area), reflecting behavior and landscape transitions, respectively. The emerged alternative use network is further analyzed through three indices: adopter large connected component (LCC) fraction, adopter assortativity, and frontier ratio of non-adopters. Adopter LCC fraction measures the degree to which adopters form a single connected component. Adopter assortativity captures the tendency of adopters to connect preferentially with one another, reflecting internal cohesion within the adopter subgraph. Frontier ratio measures the proportion of non-adopters who have at least one adopter neighbor, reflecting the breadth of active diffusion pressure across the network.

Social learning channel dynamics. IK and weighted social rewards are used to reveal social learning dynamics within different network architectures. Specifically, IK is decomposed into initial values, social increment, and individual increment. Such decomposition allows us to examine how information exchange contributes to knowledge accumulation. Weighted social rewards capture how conformity incentives evolve over simulation. Both IK and weighted social rewards are analyzed across agent types, adoption status, and transition phase to examine the temporal dynamics of each learning channel.

Landscape emergence. Emergent landscape patterns are characterized through global Moran's I, a measure of autocorrelation that quantifies whether similar values tend to occur among connected units rather than independently of proximity. Global Moran's I is calculated under two connectivity definitions: geographic adjacency, where connections are

defined by spatial proximity among parcels, and social adjacency, where connections are defined by network ties among landowners. The geographic Moran's I captures whether adopters and alternative use patches concentrate in contiguous locations or remain spatially fragmented across the geographical landscape. The social Moran's I measures whether adoption clusters within social neighborhoods. Both are computed separately for the adoption status and the alternative land use area. The composition of adoption and alternative use areas by agent type is then used to identify the distinct contributions of different landowner types to behavior change and landscape change.

2.2.5 Simulation execution

The simulation is organized into three sequential phases: network initialization, pretraining, and formal simulation.

Network initialization. The social network is constructed prior to simulation and remains fixed throughout all subsequent phases. All N landowner agents are first placed on a $W \times L$ spatial grid, with each agent occupying a unique cell representing their land parcel. Agent positions and connections are fixed across simulations under the same network type. Agent types and characteristics are independently sampled for each replication via Latin hypercube sampling (LHS), as described below. The social network is then built using the spatially embedded Watts–Strogatz (WS) small world model, where the probability of connection decays with Euclidean distance (Eq. 22). This common generative framework allows us to produce both integrated and modular social architectures by changing the average degree k and rewiring probability p. Specific parameter settings for these two topologies are detailed in Section 3. The resulting structure serves as the fixed foundation for the propagation of information exchange and normative pressure throughout the simulation.

$$p_{i,j} = Ce^{-d_{i,j}/\alpha} \tag{22}$$

Where C is a normalization constant enforcing a fixed average degree, and α is the decay parameter.

DQN pretraining. Before the formal simulation, the DQN is pretrained to develop Q-value priors that behavioral inertia under pure traditional conditions. All agents are initialized with traditional land use, with characteristics drawn from their initialization range (detailed in section 3). Agent characteristics are sampled once via LHS to ensure adequate coverage of the agent parameter space. Since pretraining involves only traditional land use with no adoption dynamics, the Q-value priors are not sensitive to network topology. The random crowd configuration with one LHS draw is considered sufficient. One shared DQN

is trained per agent type (economy-oriented, multiobjective, and amenity-oriented) using the experience replay and network update procedure described in section 2.2.3.2, running for 500 steps with 10 mini-batch updates applied per step until Q-value estimates stabilize. The pretrained weights are then saved and loaded as the initialization point for all formal simulation runs.

Formal simulation. At the beginning of formal simulation, agent types are randomly assigned to grid positions, and individual characteristics are sampled via LHS. Early adopters are seeded by probabilistically selecting agents with high degree centrality at the prescribed scale. This approach reflects the critical role of central actors in promoting information diffusion and transition (Bodin and Crona 2009). Each scenario is replicated 100 times via Monte Carlo simulation to account for stochastic variation in agent initialization and learning dynamics. Results are then aggregated for analysis. A warm-up phase is introduced before formal simulation to produce initial values for SI, $DI_{\text{performance}}$, and DI_{behavior}. During the warm-up period, DQN is active but restricted to a fixed repetition strategy with a constant ε, and IK remain static. Such settings allow early adopters to form an initial impression of the alternative use before active strategy selection begins. Following the warm-up, formal simulation begins and agents execute the "observation-strategy-decision" sequence. During formal simulation, ε decays exponentially from its warm-up endpoint. At the end of each step, 10 mini-batch updates are applied to the shared DQN for each agent type to continue learning.

3 Model parameterization and experiment design

3.1 Model parameterization

Taking forestry transition as an illustrative example, rotation forestry (RF) is designated as the traditional land use, while continuous cover forestry (CCF) serves as the alternative. Values for production service (measured by annualized timber value at a 2% discount rate), carbon sequestration, and recreation (measured by scenic beauty) are drawn from Peura et al. (2018), as presented in Table 1. Biodiversity is measured by forest-dependent species richness and expressed as relative values. Drawing on Paillet et al. (2010), the mean species richness of forest-dependent species under clear-cutting management with species compositional change is 15.79, compared to 19.95 under unmanaged forestry. This indicates that clear-cut management with species change results in approximately 20% biodiversity loss relative to an unmanaged reference condition. Accordingly, unmanaged forestry is assigned a relative biodiversity value of 1, yielding values of 0.8 and 1.0 for RF and CCF, respectively.

To ensure comparability across the four ESs, all ES values are standardized by dividing by the corresponding CCF outcome.

Table 1: Annualized ecosystem services per hectare.

	Raw Value		Standardized Value		
Ecosystem Service	**RF**	**CCF**	**RF**	**CCF**	**Source**
Annualized timber value	263.70[a]	276.53[a]	0.95	1.00	(Peura et al. 2018)
Carbon sequestration	0.23[b]	0.68[b]	0.34	1.00	(Peura et al. 2018)
Biodiversity	0.80	1.00	0.80	1.00	(Paillet et al. 2010)
Scenic beauty	5.50	6.10	0.90	1.00	(Peura et al. 2018)

[a] Unit: € year^{-1}. [b] Unit: t C year^{-1}.

Landowner agents are categorized into three types based on the forest owner behavior framework developed by Sotirov, Sallnäs, and Eriksson (2019): Economic-oriented (Optimizers), Multiobjective, and Amenity-oriented (aggregating passive, traditionalist, and environmentalist types). Agent preferences are set in two stages. First, we assign relative importance (low, medium, and high) to social rewards and four ESs, as shown in Table 2. Second, we define value ranges for social reward weights (φ) and importance scores for each ES. For each agent, a raw importance score (I_e) is randomly sampled from the corresponding range in Table 2 and normalized via Eq. (23). This normalization procedure reflects the bounded rationality of agents, capturing the inherent trade-offs between competing management objectives while maintaining the unique value orientation of each owner type. The population proportion of each landowner agent type is set according to Sotirov, Sallnäs, and Eriksson (2019).

Table 2: Preference profiles and population proportion of landowner agents.

Parameters	**Economy-oriented**	**Multiobjective**	**Amenity-oriented**
Social reward weight (φ)	Low (0, 0.1)	High (0.7, 0.9)	Low to medium (0.1, 0.6)
Provision weight (w_1)	High (8-10)	Medium (4-6)	Low (0-2)
Carbon weight (w_2)	Low (0-2)	Medium (4-6)	Medium to high (5-8)
Biodiversity weight (w_3)	Low (0-2)	Medium (4-6)	Medium to high (5-8)
Recreation weight (w_4)	Low (0-2)	Medium (4-6)	High (8-10)
Proportion (%)	35%	35%	30%

$$w_e = \frac{I_e}{\sum_{e=1}^{4} I_e} \tag{23}$$

The formal simulation runs for 150 steps. Because forestry transitions unfold over decades and may span more than one ownership period, each agent is assumed to represent a persistent ownership unit throughout the simulation. Generational turnover and land succession are not considered for simplicity. Parcel area is set to 11 hectares, which is the median declared productive forest area among natural persons (SFA, 2024). Biophysical productivity (BP) is drawn from Swedish National Forest Inventory site productivity data (2022-2024), which spans 3.2 to 8.8 $\mathrm{m^3\,ha^{-1}\,yr^{-1}}$, and is standardized to [0.36, 1.0] by dividing by the maximum value 8.8. Therefore, property similarity ($PS_{i,j}$) ranges from 0.36 to 1. CCF is still in the niche phase in Sweden, with only 3% of productive forests managed under CCF (SFA, 2024). Therefore, agents' initial implementation knowledge (IK) is randomly drawn from [0, 0.4]. Early adopters are assumed to have more CCF information and are assigned a slightly higher initial range of [0.2, 0.5]. The land share allocated to CCF by early adopters is randomly drawn from [10%, 30%] at the beginning. To ensure comparability with private utility (which ranges from 3.96 to 11), the maximum social reward (κ) is set to 7. We assume that landowner agents possess general forestry experience and can still realize some benefits (e.g., through thinning or insurance) from CCF even when lacking professional CCF knowledge. Consequently, the maximum estimation/implementation bias (B_0) is set to 0.7, which means agents can obtain at least 30% of potential ES outputs under CCF. All model parameters are listed in Table A.1.

All DQN hyperparameters are reported in Table A.2. A small strategy switching cost ($c = 0.05$) is applied to non-repetition strategies to avoid unnecessary strategy shifts. This cost stabilizes decision-making, particularly when multiple actions yield identical Q-values.

3.2 Experiment design

The simulation experiment is organized around two factors: network architecture and seeding scale. For network architecture, the Random crowd topology is configured with average degree $k = 9$ and rewiring probability $p = 0.1$. $k = 9$ slightly exceeds the Moore neighborhood, ensuring that each agent maintains social connections beyond immediate spatial adjacency and is thus exposed to a wide set of referents. The moderate p reduces global path length while maintaining small-world characteristics. The Island world topology uses $k = 3$ and $p = 0.01$. $k = 3$ falls slightly below the von Neumann neighborhood, restricting each agent's social referents to a narrow set. The near-zero p preserves the ring-lattice structure, maintaining high local clustering coefficients.

The scale of early adopters (i.e., seeding scale) is varied from 2% to 10%. Statistics from the Swedish Forestry Agency show that only 3% of productive forests currently apply

CCF, which may be even lower among family-owned forestry. The Swedish government has committed to achieving 20% CCF under the EU 2030 Forest Strategy. The seeding scale range spans the empirically plausible near-term trajectory of the transition. Each scenario is replicated 100 times via Monte Carlo simulation to account for stochastic variation in agent initialization and learning dynamics. Simulation experiment settings are presented in Table 3.

Table 3: Simulation experiments.

Experiment	**Network architecture**	**Seeding scale**
1	Random crowd ($k = 9$, $p = 0.1$)	$[2\%, 10\%]$
2	Island world ($k = 3$, $p = 0.01$)	$[2\%, 10\%]$

4 Results

Results follow the analytical structure of the conceptual framework. Section 4.1 characterizes aggregate transition outcomes across network architectures, establishing how social structure shapes the formation of the alternative use network. Section 4.2 opens the social learning engine to examine how information exchange and normative pressure operate across transition phases, adoption status, and landowner types, and how network architecture modulates each channel. Section 4.3 analyzes the landscape patterns that emerge from these dynamics, identifying which landowner types drive geographic and social land use change under different architectures.

4.1 Transition outcomes under different network architecture

Figure 3 compares final adoption outcomes of different network structures and varying seeding scales. Random crowd yields higher adoption rates and greater alternative land use area than island world across all seeding scales, implying that connectivity plays a significant role in sustaining alternative use network.

A significant gap exists between the adoption rate and the actual alternative use area, particularly as the seeding scale increases. This disparity may reflect two features of land use transition. First, land use transition may proceed through many small-scale experiments distributed across a large number of landowners. Second, it may reflect the realistic medium-term reach of alternative practices: widespread in acceptance but modest in implementation depth.

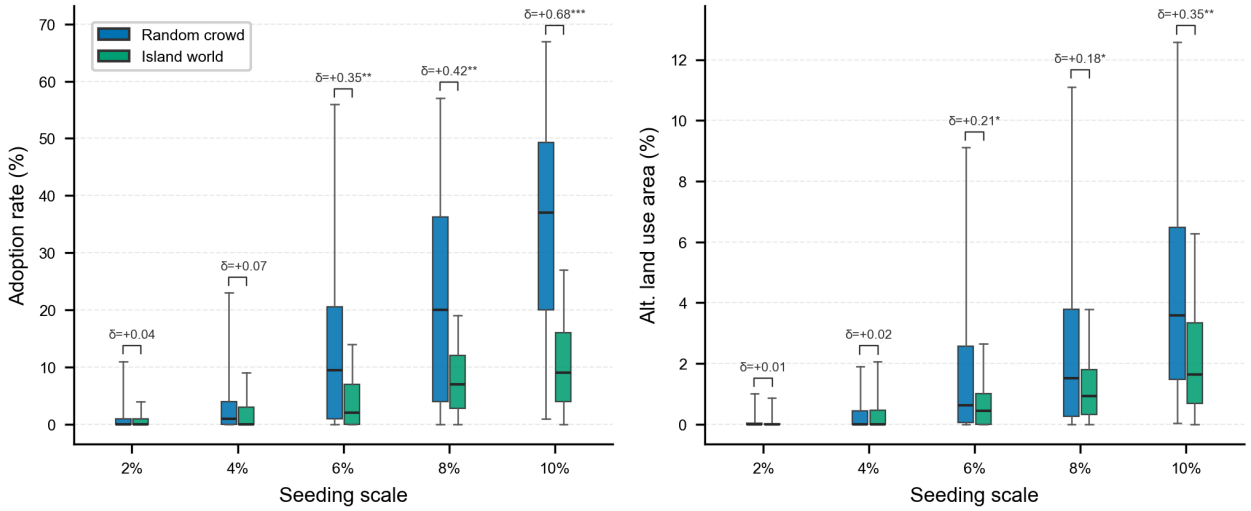

Figure 3: Final adoption and alternative land use under varying seeding scales across network structures. δ denotes Cliff's delta for the comparison between random crowd and island world; * denotes significant differences.

Figure 4 shows three indicators of how the alternative use network evolves over time. Random crowd produces rapid consolidation and broad diffusion pressure. Large connected components (LCC) fraction reaches 1.0 within the first 30 steps, reflecting that all adopters quickly form one connected group. Frontier ratio remains substantially higher than in the island world throughout the simulation, indicating sustained and broad contact between adopters and non-adopters.

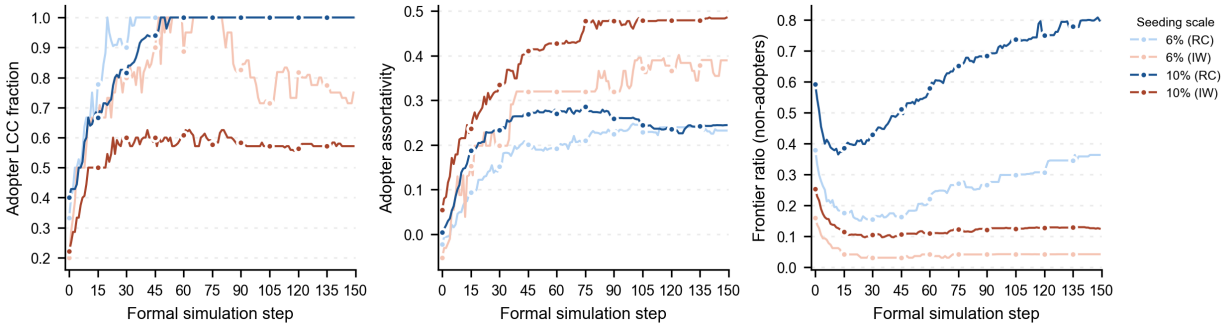

Figure 4: Connectivity and diffusion reach within the alternative use network across seeding scales and topologies. **LCC**, large connected components; **RC**, random crowd; **IW**, island world).

In contrast, island world shows fragmented adoption clusters but higher internal cohesion. Final LCC fraction in island world is significantly lower than in random crowd, meaning higher fragmentation. Notably, LCC fraction is lower at 10% seeding scale than at 6% seeding scale in island world. This implies that higher seeding scales may activate hubs

across multiple clusters simultaneously, but sparse cross-group ties prevent these clusters from merging. Consequently, more early adopters paradoxically fragment the alternative use network. However, adopter assortativity is consistently higher in island world, indicating that adopters within each local group are tightly interconnected.

These structural differences explain the transition gap between network architectures. Random crowd sustains broad diffusion pressure through high frontier exposure. Island world concentrates normative pressure within local groups through high assortativity, but limits its reach across the network through low frontier ratio and fragmented connectivity. Therefore, modularity can not only impede information exchange but also create higher normative pressure, thus explaining worse transition results in island world.

4.2 Role of social learning engine

Figure A.1 shows the temporal evolution of adoption rate and alternative land use area across seeding scales and network structures. The trajectories reveal three distinct phases: an initial reversion (steps 0–15) in which adoption rate and alternative land use area decline, a middle phase (steps 16–120) in which low seeding scales stagnate while high seeding scales grow, and a final stabilization phase (steps 121–150). Based on these trajectories, steps 0–15, 50–70, and 130–150 are selected as representative periods for each phase, with steps 10, 60, and 140 as representative time slices. Seeding scale 10% is used as the representative case for subsequent analysis.

Figure 5 decomposes implementation knowledge (IK) into initial values, social increments, and individual increments at three representative time slices. It shows that, individual learning dominates IK accumulation for adopters throughout the simulation. Social learning gradually becomes an important way of knowledge acquisition in later stages, but remains secondary to individual learning in both network structures.

Network structure can affect the social learning channel. In random crowd, social IK increment is larger than in island world across all agent types, confirming that modular network architecture impedes information exchange. This effect is more pronounced for economy-oriented agents: in random crowd, social learning accounts for a substantial share of their IK at the final phase, whereas in island world its contribution remains limited, leaving economy-oriented agents with lower overall IK under the modular topology.

Non-adopters accumulate IK almost exclusively through social learning, yet remain significantly lower than adopters throughout the simulation. This divergence demonstrates that social exposure alone is insufficient to close the IK gap. Without the experiential reinforcement that follows adoption, non-adopters receive signals but cannot convert them into the

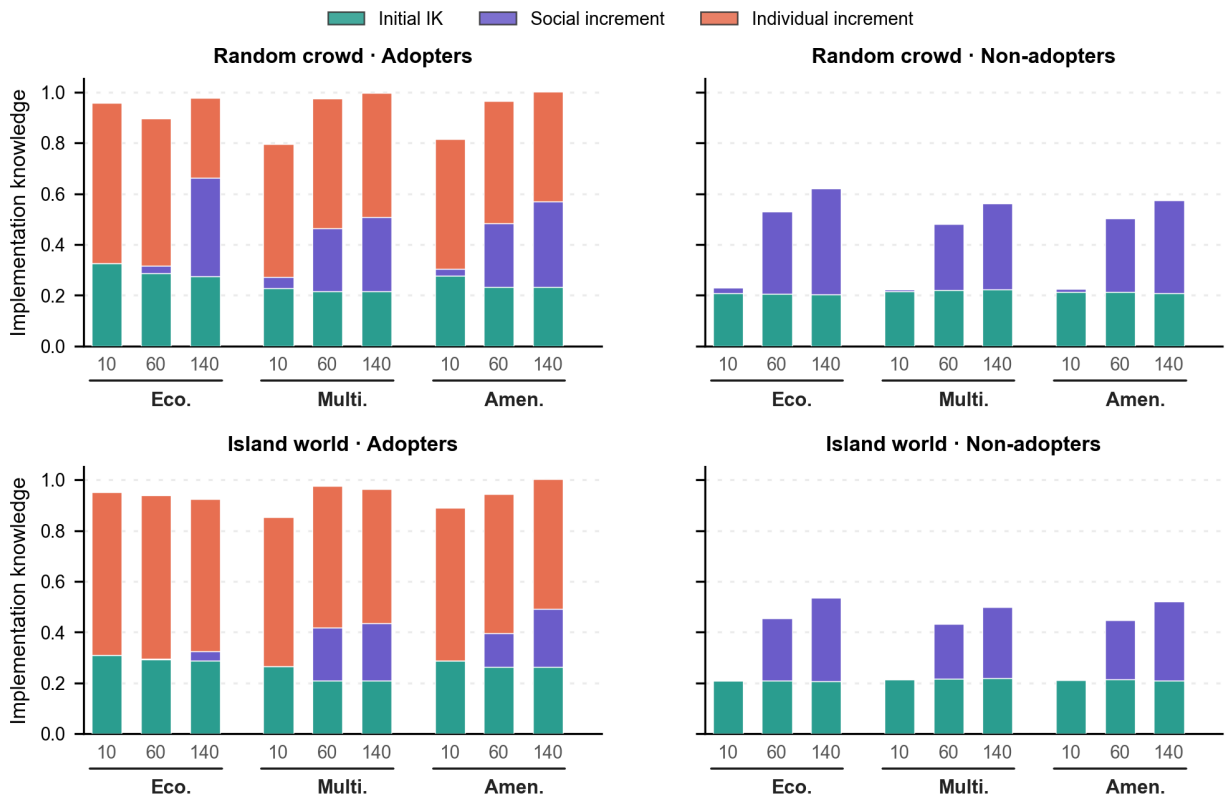

Figure 5: Decomposition of implementation knowledge IK components prior to adoption events across transition phases under seeding scale = 10%. IK levels are categorized into initial values, social increments (social learning), and individual increments (individual learning) for both adopters and non-adopters; the x-axis within each agent type group shows representative time slices (steps 10, 60, 140).

knowledge levels required for transition.

Figure 6 shows weighted social rewards at representative time slices across agent types, adoption status, and network architectures. Since social reward is maximized when an agent's land use practices (i.e., alternative use ratio) match its neighborhood average, high values indicate strong conformity with the prevailing local norm.

Non-adopters generally receive higher social rewards than adopters throughout the simulation, reflecting sustained conformity with a predominantly traditional-use neighborhood. Given that overall alternative use ratios remain low, most non-adopters are rarely surrounded by a majority of intensive alternative users, keeping the conformity benefit of remaining traditional use intact. This creates a lock-in effect: switching to alternative use would reduce social rewards.

Notably, amenity-oriented adopters in random crowd receive social rewards comparable to or higher than non-adopters. This implies that their neighborhoods have shifted sufficiently

toward alternative use to sustain conformity.

Adopters show increasing social rewards over time in random crowd, consistent with the growing alternative use network. In island world, this dynamic reverses: economy-oriented and amenity-oriented adopters' social rewards decline over time. This may appear counterintuitive given that island world produces higher adopter assortativity. However, social rewards reflect conformity with the neighborhood's prevailing implementation intensity, not simply whether neighbors have adopted. Even within tight adopter clusters, the local norm remains anchored at low implementation levels. Agents who implement more deeply than their neighbors increasingly deviate from this low-implementation local norm over time, reducing their social rewards. Therefore, the modular structure amplifies the local norm of low implementation intensity. This explains why adoption cohesion and implementation lock-in coexist in island world.

This mechanism also explains the adoption-landscape gap observed in Figure 3. Normative pressure does not primarily suppress adoption decisions, as private utility from alternative use exceeds traditional use across all agent types. Instead, conformity pressure constrains how much landowners implement, anchoring implementation to the prevailing low neighborhood average. High acceptance therefore coexists with low implementation depth because deviating upward from the local norm incurs a social penalty that offsets the private incentive for deeper implementation. This suggests that voluntary adoption alone is insufficient to generate enough normative pressure to overcome traditional use lock-in, even when alternative use can yield higher private utility.

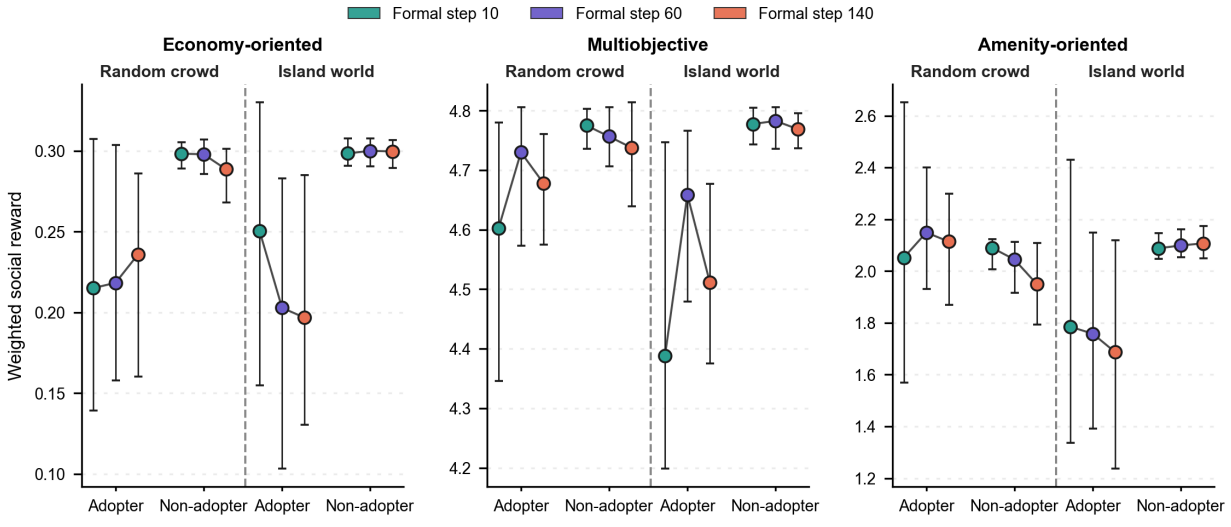

Figure 6: Weighed social rewards at representative time steps under seeding scale = 10%. Social rewards reach their highest when the agent's alternative use ratio matches the local group average.

4.3 Landscape pattern

Figure 7 shows global Moran's I of alternative use at the simulation end. This index is computed separately for geographic and social adjacency, for both adoption binary status and alternative use ratio.

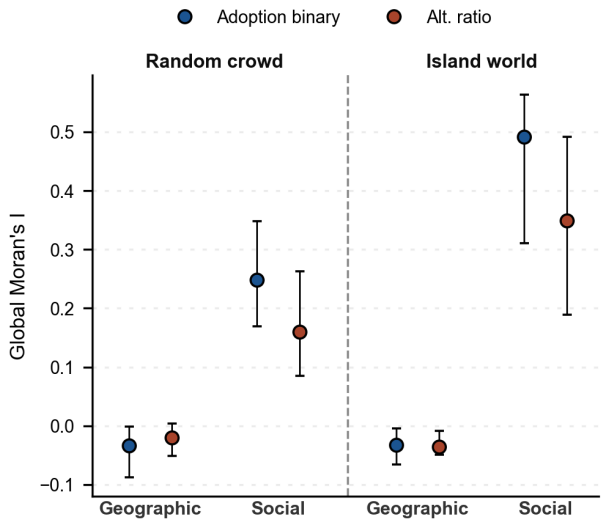

Figure 7: Global Moran's I of alternative land use computed under geographic adjacency and social adjacency at the end of the simulation. Moran's I is computed separately for a binary adoption status and alternative land use ratio. Results shown for seeding scale=10%.

The results show that alternative use is socially contagious along network ties, not geographical adjacency. Social Moran's I for adoption status reaches 0.26 in random crowd and 0.49 in island world, while geographic Moran's I remains near zero across both network types. Although both networks embed distance decay in tie formation, adoption does not cluster geographically. The geographic dissociation arises because agent characteristics are spatially randomized, indicating that the social architecture of landowner interactions, not their spatial arrangement, is the primary determinant of transition patterns.

Social Moran's I under random crowd is approximately half that of island world for both adoption binary status and alternative use ratio. This indicates that adopters cluster within local groups and create cohesive blocks in island world, while adopters spread through long-range ties and produce a more diverse mosaic pattern in random crowd.

Binary Moran's I exceeds alternative use ratio Moran's I by approximately 0.09 in random crowd and 0.15 in island world, indicating that neighbors converge more strongly on whether to adopt than on how much to implement. The larger gap in island world reflects the coexistence of two dynamics: tight clustering causes neighbors to become structurally more similar to one another, while prevailing implementation norm within these clusters remains

low. Agents who implement more deeply than their neighbors deviate upward from the shallow local norm. This further confirms that the normative pressure locks the implementation intensity at a low level.

Figure 8 and Figure A.2 show adoption rate and alternative land use area by agent type at three representative time steps within random crowd and island world, respectively. The results show that landowner types play distinct roles in adoption and implementation, and these roles shift as network type and seeding scale change.

At low seeding scale, network architecture affects which landowner types remain active. In random crowd at 6% seeding scale, all three types have comparable adoption rates throughout the simulation, with economy-oriented agents contributing the most to the alternative use landscape. In island world at the same seeding scale, multiobjective agents drop to near-zero adoption by mid-simulation, and amenity-oriented agents contribute negligibly to landscape change despite maintaining some adoption activity. This contrast implies that the modular structure of island world impose great normative pressure of transitional use on landowners with higher social conformity. This makes economy-oriented agents the major contributors to both behavior and landscape transition in a modular society at the niche phase.

High seeding scale can create a favorable social environment and activate amenity-oriented landowners regardless of network type. At 10% seeding scale, amenity-oriented agents have both a high adoption rate and alternative use ratio throughout simulation regardless of network type. This effect is stronger in random crowd, where integrated connectivity allows new norm of alternative use to propagate more broadly.

Economy-oriented agents have high implementation intensity across network types and seeding scales. This implies that they are the most robust contributors to landscape change despite accounting for the least contribution to adoption rate.

Multiobjective agents are important for behavior change while show the largest gap between adoption rate and implementation intensity. Especially when seeding scale = 10%, they adopt at high rates but contribute least to alternative use landscape. This suggests that adoption decisions and implementation intensity are driven by different mechanisms: social conformity may trigger adoption, but individual attributes shape how the practice is implemented.

5 Discussion

This study examined how two social learning channels (information exchange and normative pressure) drive land use transition under contrasting network architectures. The following three aspects are discussed in this section: the distinct and phase-sequential roles of each

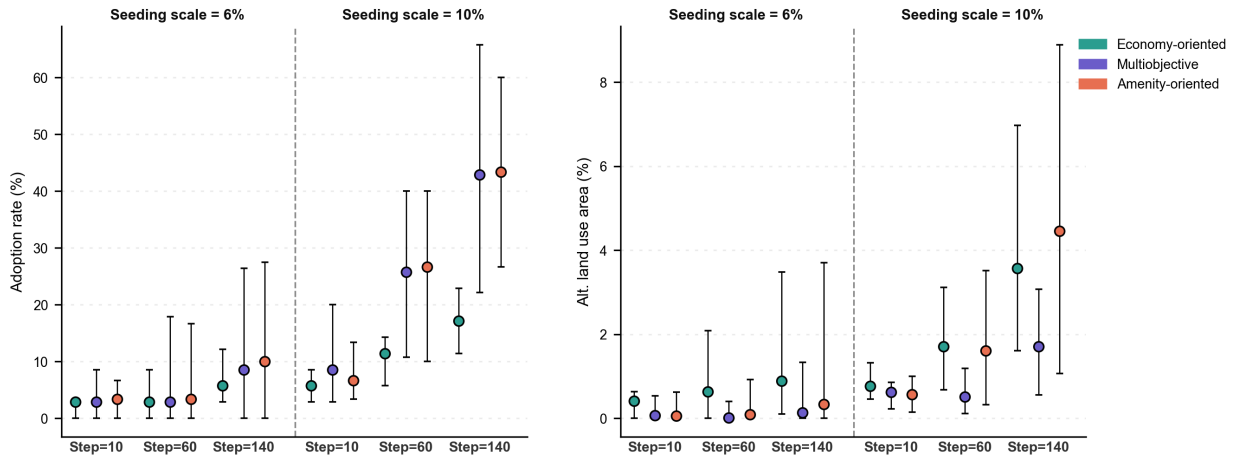

Figure 8: Adoption and alternative land use by agent type at representative time slices within random crowd.

learning channel, how network architecture shapes these channels, and how the resulting landscape patterns diverge depending on both architecture and landowner composition.

5.1 Roles of information exchange and normative pressure

The simulation results reveal that information exchange shapes the feasibility of adopting alternative use, while normative pressure determines how deeply alternative practice is implemented.

In early transition, information exchange through social learning is limited because population-wide implementation knowledge (IK) is too low to circulate meaningful knowledge. Therefore, IK accumulates primarily through individual practice in the beginning. As transition progresses, social learning becomes increasingly effective and begins to narrow the knowledge gap in mid-to-late phases, though it cannot fully substitute for individual learning. This limitation is most acute for non-adopters, who depend on peer information yet remain knowledge-constrained throughout the transition.

Normative pressure does not block adoption but constrains implementation depth. Since alternative use yields higher private utility across all landowner types, adoption is privately attractive regardless of social architecture. What normative pressure does instead is anchor implementation intensity to the prevailing neighborhood average. When most neighbors practice traditional use or low-level alternative use, increasing the alternative use ratio beyond the local norm incurs a social utility reduction. This mechanism explains the gap between the adoption rate and the actual alternative use area. Landowners accept the practice but implement it shallowly to avoid social costs. Normative pressure is therefore biased

toward whichever norm is dominant. Since the transition begins from a predominantly traditional landscape, it systematically constrains implementation depth.

5.2 Effects of network architecture on learning channels

Network architecture shapes the information exchange channel differently depending on landowner type. In random crowd, broad connectivity gives agents wider access to practicing adopters, producing larger social IK increments. This benefit is strongest for economy-oriented agents, whose transition is most sensitive to knowledge improvement. In island world, sparse cross-community ties restrict this channel. Economy-oriented agents consequently accumulate lower IK, slowing their transition. Network structure thus determines how effectively social ties convert into usable knowledge.

The effect on normative pressure reveals a separation between network structure and land use outcomes. Island world produces high adopter assortativity: adopters connect preferentially with other adopters. Yet this social clustering does not reinforce implementation intensity and enhance the social rewards of alternative uses. The reason is that assortativity describes who connects with whom, while normative pressure governs alignment in practice intensity. When the prevailing local norm is anchored at low implementation, tight clustering sustains that shallow norm rather than driving intensive implementation. Social cohesion and implementation lock-in therefore coexist in island world. In random crowd, by contrast, the growing adopter network gradually raises the neighborhood average, enhancing the social rewards of alternative use over time.

5.3 Social landscape configuration and landowner type contributions

Global Moran's I indicates that landscape change diffuse through social ties rather than geographic proximity. Adoption is organized by who landowners are socially connected to, not by where their parcels are located, even though geographic distance decay is considered in network generation. Network architecture further shapes how adoption clusters within this social network. Island world generates cohesive adoption blocks in the network, where social neighborhoods converge on shared adoption and implementation norms, forming locally intensive but isolated clusters among connected landowners. Random crowd produces a diffuse social mosaic, with adopters distributed broadly across long-range ties without forming dense local clusters. These patterns are reflected by Moran's I results: social Moran's I in island world is higher than in random crowd, while geographic Moran's I remains near zero in both architectures. This suggests that spatially targeted policy instruments may

have limited reach in steering transition diffusion, and that interventions designed around social network position are more likely to align with the actual organizing logic of landscape change.

Within both social architectures, neighbors converge on whether to adopt but diverge on how much to implement. Social learning accelerates adoption by circulating implementation knowledge through network ties, lowering the knowledge barrier for potential adopters. However, implementation intensity remains governed by individual preferences and conformity to the local norm.

This divergence is reflected in the landowner type contributions. Economy-oriented agents are weakly sensitive to social conformity and therefore implement primarily according to private returns, contributing most to actual landscape change. Multiobjective agents are highly conformist. They adopt widely but implement minimally, shifting social norms more than changing the land. Amenity-oriented agents are conditionally activated. Once the local norm tips toward alternative use, social and private incentives align, enabling deeper implementation. These different response patterns imply that encouraging adoption and driving landscape change require different levers, and instruments targeting social acceptance may need to be complemented by those directly engaging economy-oriented landowners.

5.4 Implications

Our findings suggest that social dynamics and network architecture play a significant role in transition governance. In early phases, the primary barrier is knowledge access rather than social resistance, and interventions that support knowledge accumulation through technical assistance are most relevant. As adoption spreads and population-wide knowledge rises, the binding constraint shifts from information access to conformity pressure. At this stage, interventions that raise the visible norm of implementation intensity become more effective than those targeting further adoption acceptance.

Network architecture shapes which outreach strategies are likely to work. In integrated networks, broad information circulation can support adoption diffusion by widening access to knowledgeable adopters. In modular networks, where tight clustering sustains existing norms, consolidating adoption and implementation norms within communities before extending outreach may be more effective than wide or diffuse seeding.

Our results suggest that landowner heterogeneity shapes which instruments are likely to produce landscape change versus behavior change. Landowners whose decisions are primarily driven by private utility contribute most to actual landscape change and respond most directly to knowledge improvement. Those who are more sensitive to social norms tend

to support broader adoption but contribute less to implementation intensity. Those who respond conditionally to local norms are better reached through norm consolidation than direct outreach. Because adoption rate and landscape change reflect distinct social processes, instruments designed to increase behavior change need to be paired with those that directly support implementation intensity among landowners whose decisions are less conformity-constrained.

6 Conclusions

Successful land use transition depends on social coordination and learning processes embedded in network structures. This study developed SALT (**S**ocial learning in **A**gent-based **L**and use **T**ransitions), a spatially explicit agent-based model, to examine how social learning interacts with two contrasting network architectures to shape land use transition. The model is parameterized using a Swedish forestry context. Based on simulation results, we analyze two social learning channels, information exchange and normative pressure, and explore their dynamics across integrated and modular networks.

Our results show that the two social learning channels play distinct roles that shift across transition phases. In early transition, implementation knowledge is the binding constraint on adoption feasibility, and individual experiential learning is the primary driver of knowledge accumulation. Social learning narrows the knowledge gap in later phases but cannot fully substitute for individual learning. This suggest that, landowners need to be active practitioners of alternative use, not only recipients of social information, to push the transition forward. The common observation that many landowners express interest in alternative practices, yet transition remains slow may reflect not a lack of motivation but a persistent knowledge gap that social networks alone cannot close. As adoption spreads and knowledge becomes less limiting, normative pressure emerges as the dominant constraint, anchoring implementation intensity to the prevailing local average and explaining the gap between adoption rate and actual landscape change.

Network architecture shapes both channels but in different ways. Integrated networks widen information access and allow the alternative-use norm to strengthen as the adopter community grows. Modular networks restrict information circulation and lock in low-implementation norms locally, and higher seeding can paradoxically fragment the alternative use network by activating clusters that lack cross-group ties to merge.

Landscape change organizes along social ties rather than geographic proximity, with architecture determining whether adoption forms cohesive blocks or a diffuse mosaic in social network. Across both architectures, neighbors converge on whether to adopt but diverge on

how much to implement, and landowner types contribute differently to behavior change and landscape change. Effective transition governance must therefore consider both transition phases and social dynamics, as these together determine which learning channel is binding, which landowner types are reachable, and what landscape outcomes emerge.

Our results also suggests what land use transitions may look like in contexts starting from high conformity. Even where alternative practices generate higher private returns, transition is likely to result in a relatively large share of landowners accepting the alternative practice while only a minority of the land area actually transitions. This study mainly focused on the social learning engine and how its two channels operate across transition phases and network architectures. The extent to which such a transition produces a genuinely multifunctional landscape, and whether the ecological thresholds required for meaningful ecosystem service delivery are reached, remain important questions for future research.

References

Blattert, C. et al. (2023). "Climate targets in European timber-producing countries conflict with goals on forest ecosystem services and biodiversity". *Communications Earth & Environment* 4.1, p. 119. DOI: 10.1038/s43247-023-00771-z.

Bodin, Ö. and B. I. Crona (2009). "The role of social networks in natural resource governance: What relational patterns make a difference?" *Global Environmental Change* 19.3, pp. 366–374. DOI: 10.1016/j.gloenvcha.2009.05.002.

Deutsch, M. and H. B. Gerard (1955). "A study of normative and informational social influences upon individual judgment". *The Journal of Abnormal and Social Psychology* 51.3, pp. 629–636. DOI: 10.1037/h0046408.

Dong, Q. et al. (2022). "An adaptive group decision making framework: Individual and local world opinion based opinion dynamics". *Information Fusion* 78, pp. 218–231. DOI: 10.1016/j.inffus.2021.09.013.

Duinen, R. van et al. (2016). "Going beyond perfect rationality: drought risk, economic choices and the influence of social networks". *The Annals of Regional Science* 57.2, pp. 335–369. DOI: 10.1007/s00168-015-0699-4.

Eastwood, A. et al. (2022). "A cup of tea? – The role of social relationships, networks and learning in land managers' adaptations to policy change". *Land Use Policy* 113, p. 105926. DOI: 10.1016/j.landusepol.2021.105926.

Efferson, C., S. Vogt, and E. Fehr (2020). "The promise and the peril of using social influence to reverse harmful traditions". *Nature Human Behaviour* 4.1, pp. 55–68. DOI: 10.1038/s41562-019-0768-2.

Emami-Skardi, M. J., N. Momenzadeh, and R. Kerachian (2021). "Social learning diffusion and influential stakeholders identification in socio-hydrological environments". *Journal of Hydrology* 599, p. 126337. DOI: 10.1016/j.jhydrol.2021.126337.

Evans, D. et al. (2024). "Carbon farming diffusion in Australia". *Global Environmental Change* 89, p. 102921. DOI: 10.1016/j.gloenvcha.2024.102921.

Gu, J. (2024). "Neighborhood Does Matter: Farmers' Local Social Interactions and Land Rental Behaviors in China". *Land* 13.1, p. 76. DOI: 10.3390/land13010076.

Hertog, I. M., S. Brogaard, and T. Krause (2022). "Barriers to expanding continuous cover forestry in Sweden for delivering multiple ecosystem services". *Ecosystem Services* 53, p. 101392. DOI: 10.1016/j.ecoser.2021.101392.

Huber, R. et al. (2022). "Bridging behavioural factors and standard bio-economic modelling in an agent-based modelling framework". *Journal of Agricultural Economics* 73.1, pp. 35–63. DOI: 10.1111/1477-9552.12447.

Hung, F. and Y. C. E. Yang (2021). "Assessing Adaptive Irrigation Impacts on Water Scarcity in Nonstationary Environments—A Multi-Agent Reinforcement Learning Approach". *Water Resources Research* 57.9, e2020WR029262. DOI: 10.1029/2020WR029262.

IPCC (2019). *Climate Change and Land: an IPCC special report on climate change, desertification, land degradation, sustainable land management, food security, and greenhouse gas fluxes in terrestrial ecosystems*. Cambridge, UK and New York, NY, USA: Cambridge University Press. 896 pp.

Isaac, M. and P. Matous (2017). "Social network ties predict land use diversity and land use change: a case study in Ghana". *Regional Environmental Change* 17.6, pp. 1823–1833. DOI: 10.1007/s10113-017-1151-3.

Jager, W. et al. (2000). "Behaviour in commons dilemmas: *Homo economicus* and *Homo psychologicus* in an ecological-economic model". *Ecological Economics* 35.3, pp. 357–379. DOI: 10.1016/S0921-8009(00)00220-2.

Janssen, M., C. Vlek, and W. Jager (1999). *Consumats in a commons dilemma: Testing the behavioral rules of simulated consumers*. Groningen: RUG-COV.

Junquera, V. and A. Grêt-Regamey (2019). "Crop booms at the forest frontier: Triggers, reinforcing dynamics, and the diffusion of knowledge and norms". *Global Environmental Change* 57, p. 101929. DOI: 10.1016/j.gloenvcha.2019.101929.

Kreft, C. et al. (2024). "Quantifying the impact of farmers' social networks on the effectiveness of climate change mitigation policies in agriculture". *Journal of Agricultural Economics* 75.1, pp. 298–322. DOI: 10.1111/1477-9552.12557.

Kremen, C. and A. M. Merenlender (2018). "Landscapes that work for biodiversity and people". *Science* 362.6412, eaau6020. DOI: 10.1126/science.aau6020.

Loch, T. K. and D. Kleinschmit (2025). "Building connections: Exploring social network research in forest sciences". *Forest Policy and Economics* 170, p. 103382. DOI: 10.1016/j.forpol.2024.103382.

Luján Soto, R. et al. (2021). "Participatory monitoring and evaluation to enable social learning, adoption, and out-scaling of regenerative agriculture". *Ecology and Society* 26.4. DOI: 10.5751/ES-12796-260429.

Manson, S. M. et al. (2016). "Modeling the effect of social networks on adoption of multifunctional agriculture". *Environmental Modelling & Software* 75, pp. 388–401. DOI: 10.1016/j.envsoft.2014.09.015.

Martin, D. A. et al. (2022). "Land-use trajectories for sustainable land system transformations: Identifying leverage points in a global biodiversity hotspot". *Proceedings of the National Academy of Sciences* 119.7, e2107747119. DOI: 10.1073/pnas.2107747119.

Matous, P. and Ö. Bodin (2024). "Hub-and-spoke social networks among Indonesian cocoa farmers homogenise farming practices". *People and Nature* 6.2, pp. 598–609. DOI: 10.1002/pan3.10578.

Muthukrishna, M. and M. Schaller (2020). "Are Collectivistic Cultures More Prone to Rapid Transformation? Computational Models of Cross-Cultural Differences, Social Network Structure, Dynamic Social Influence, and Cultural Change". *Personality and Social Psychology Review: An Official Journal of the Society for Personality and Social Psychology, Inc* 24.2, pp. 103–120. DOI: 10.1177/1088868319855783.

Paillet, Y. et al. (2010). "Biodiversity Differences between Managed and Unmanaged Forests: Meta-Analysis of Species Richness in Europe". *Conservation Biology* 24.1, pp. 101–112. DOI: 10.1111/j.1523-1739.2009.01399.x.

Peura, M. et al. (2018). "Continuous cover forestry is a cost-efficient tool to increase multifunctionality of boreal production forests in Fennoscandia". *Biological Conservation* 217, pp. 104–112. DOI: 10.1016/j.biocon.2017.10.018.

Phelps, C., R. Heidl, and A. Wadhwa (2012). "Knowledge, Networks, and Knowledge Networks: A Review and Research Agenda". *Journal of Management* 38.4, pp. 1115–1166. DOI: 10.1177/0149206311432640.

Sanderson, D. R., T. P. McAllister, and J. Helgeson (2025). "Simulating future household adaptation to sea level rise using agent-based modeling and reinforcement learning". *International Journal of Disaster Risk Reduction* 128, p. 105742. DOI: 10.1016/j.ijdrr.2025.105742.

Skaalsveen, K., J. Ingram, and J. Urquhart (2020). "The role of farmers' social networks in the implementation of no-till farming practices". *Agricultural Systems* 181, p. 102824. DOI: 10.1016/j.agsy.2020.102824.

Slijper, T. et al. (2022). "Exploring how social capital and learning are related to the resilience of Dutch arable farmers". *Agricultural Systems* 198, p. 103385. DOI: 10.1016/j.agsy.2022.103385.

Sotirov, M., O. Sallnäs, and L. O. Eriksson (2019). "Forest owner behavioral models, policy changes, and forest management. An agent-based framework for studying the provision of forest ecosystem goods and services at the landscape level". *Forest Policy and Economics* 103, pp. 79–89. DOI: 10.1016/j.forpol.2017.10.015.

Spencer, G. M. (2012). "Creative economies of scale: an agent-based model of creativity and agglomeration". *Journal of Economic Geography* 12.1, pp. 247–271. DOI: 10.1093/jeg/lbr002.

Stryamets, N. et al. (2026). "To clear-cut or not to clear-cut: Diversifying benefits from small-scale forestry in Sweden". *Forest Ecosystems* 15, p. 100401. DOI: 10.1016/j.fecs.2025.100401.

Terasaki Hart, D. E. et al. (2023). "Priority science can accelerate agroforestry as a natural climate solution". *Nature Climate Change* 13.11, pp. 1179–1190. DOI: 10.1038/s41558-023-01810-5.

Valujeva, K. et al. (2023). "Pathways for governance opportunities: Social network analysis to create targeted and effective policies for agricultural and environmental development". *Journal of Environmental Management* 325, p. 116563. DOI: 10.1016/j.jenvman.2022.116563.

Wang, Y. et al. (2021). "Role of social networks in building household livelihood resilience under payments for ecosystem services programs in a poor rural community in China". *Journal of Rural Studies* 86, pp. 208–225. DOI: 10.1016/j.jrurstud.2021.05.017.

Wolf, I. et al. (2015). "Changing minds about electric cars: An empirically grounded agent-based modeling approach". *Technological Forecasting and Social Change* 94, pp. 269–285. DOI: 10.1016/j.techfore.2014.10.010.

Zhang, W., A. Valencia, and N.-B. Chang (2023). "Synergistic Integration Between Machine Learning and Agent-Based Modeling: A Multidisciplinary Review". *IEEE Transactions on Neural Networks and Learning Systems* 34.5, pp. 2170–2190. DOI: 10.1109/TNNLS.2021.3106777.

A Appendix

A.1 Model parameter and DQN setting

Table A.1: Model parameters.

Parameter	**Type**	**Value (Unit)**	**Description**
$area$	Exogenous	11 (ha)	Land area of the parcel.
BP_i	Exogenous	[0.36, 1.0] (scalar)	Biophysical productivity, reflecting variation in site quality.
$UES_{e,i}^{\text{trad/alt}}$	Exogenous	Tabel 1	e-th annualized ES outputs per unit area of traditional ("trad") or alternative use ("alt"). $e = 1$: provision; $e = 2$: carbon stock; $e = 3$: biodiversity; $e = 4$: recreation.
$UES_{e,i,t}$	Endogenous	-	e-th annualized ES outputs per unit area of parcel i.
$w_{e,i}$	Exogenous	Tabel 2	Preference weight of e-th ES.
φ_i	Exogenous	Tabel 2	Weight of social rewards.
IK	Exogenous (initial)	Early adopters: [0.2, 0.5]; Others:[0, 0.4] (scalar)	Implementation knowledge.
ML_i	Exogenous	10 (year)	Memory length.
sw	Exogenous	30 (%)	Searching width for bounded innovation under the optimization strategy.
l^*	Exogenous	[0.3, 0.7] (scalar)	Learning rate.
l	Endogenous	Randomly drawn from $[0, l^*]$	Update rate of implementation knowledge.
$SI_{i,t}$	Endogenous	-	Satisfaction index.
$DI_{i,t}$	Endogenous	-	Dissimilarity index.
$DI_{i,t}^{\text{b}}$	Endogenous	-	Behavioral deviation.
$DI_{i,t}^{\text{p}}$	Endogenous	-	Performance disparity.
$U_{i,t}^{\text{private}}$	Endogenous	-	Utility from annualized four ES outputs.
$U_{i,t}^{\text{social}}$	Endogenous	-	Social rewards from personal network.
κ	Exogenous	7 (scalar)	Maximum social rewards.
$PA_{i,t}$	Exogenous (initial)	Early adopters: [10, 30]; Others: 0 (%)	Proportion of alternative use applied by agent i.
$\overline{PA}_{i,t}^{\text{Nei}}$	Endogenous	-	Average proportion of alternative use applied within the personal network.
N_i^{Nei}	Structural	-	Number of peers (i.e., direct connections).
$GP_{i,j}$	Endogenous	-	Geographical proximity between agents i and j.

Continued on next page

Table A.1 – continued

Parameter	Type	Value (Unit)	Description
$d_{i,j}$	Structural	-	Geographical distance between agents i and j.
$NS_{i,j}$	Structural	-	Structural similarity of agents i's and j's personal networks.
$NM_{i,j}$	Structural	-	The number of common connections.
$PS_{i,j}$	Exogenous	[0.36, 1] (scalar)	Property similarity (i.e., the similarity of parcels' biophysical productivity).
$H_{i,j}$	Endogenous	-	Homophily between agents i and j.
B_0	Exogenous	0.7 (scalar)	Maximum estimation bias of unit ES output.
ε	Endogenous	[0.8, 1.2]	Random risk.

* **Note: Exogenous**: Fixed externally. **Exogenous (initial)**: Set externally at initialization, but updated endogenously during simulation. **Endogenous**: Fully generated and updated within the model. **Structural**: Derived from model structure.

Table A.2: DQN hyperparameters.

Parameter	Value	Phase	Description
Network architecture			
State input	$[SI, DI^{\mathrm{p}}, DI^{\mathrm{b}}]$ (dim=3)	All	State vector fed to Q-network
Action output	4 strategies	All	Q-values over repetition, optimization, imitation, social comparison
Hidden layers	2×64 units, ReLU	All	Fully connected MLP
Optimizer	Adam	All	Gradient-based weight update
Replay buffer size	1,000	All	Maximum stored transitions per agent type
Pretraining			
Steps	500	Pretrain	Steps run under pure traditional use, no early adopters
LR	1×10^{-3}	Pretrain	Adam optimizer learning rate
Discount factor γ	0.98	Pretrain	Weight on future vs. immediate rewards

Continued on next page

Table A.2 – continued

Parameter	Value	Phase	Description
Initial ε	1.0	Pretrain	Full exploration at start
Minimum ε	0.2	Pretrain	Exploration floor
ε decay rate	0.995	Pretrain	Per replay call, multiplicative
Batch size	128	Pretrain	Transitions sampled per replay call
Replay calls per step	10	Pretrain	Q-network updates per simulation step
Target update interval	10 steps	Pretrain	Frequency of target network sync
Formal warmup			
Steps	5	Warmup	Fixed repetition before active strategy selection begins
ε	0.2	Warmup	Held constant throughout warmup
Replay calls per step	1	Warmup	Minimal training during warmup
Formal simulation			
Steps	150	Formal	Total steps after warmup
LR	1×10^{-4}	Formal	Reduced LR for fine-tuning on pretrained weights
Discount factor γ	0.98	Formal	Weight on future vs. immediate rewards
Initial ε	0.2	Formal	Reset at start of formal phase
Minimum ε	0.01	Formal	Exploration floor
ε decay rate	0.98	Formal	Per step, multiplicative
Batch size	32	Formal	Transitions sampled per replay call
Replay calls per step	1	Formal	One update per step
Target update interval	10 steps	Formal	Frequency of target network sync
Strategy switching cost c			
Repetition	0.0	All	No cost for staying with current strategy
Optimization	0.05	All	Cost subtracted from stored reward when switching
Imitation	0.05	All	Cost subtracted from stored reward when switching

Continued on next page

Table A.2 – continued

Parameter	Value	Phase	Description
Social comparison	0.05	All	Cost subtracted from stored reward when switching

* **Note:** LR: learning rate. ε: exploration probability in ε-greedy policy. Replay calls per step: number of `replay` calls applied to the shared DQN per agent type at the end of each simulation step.

A.2 Supplementary figure

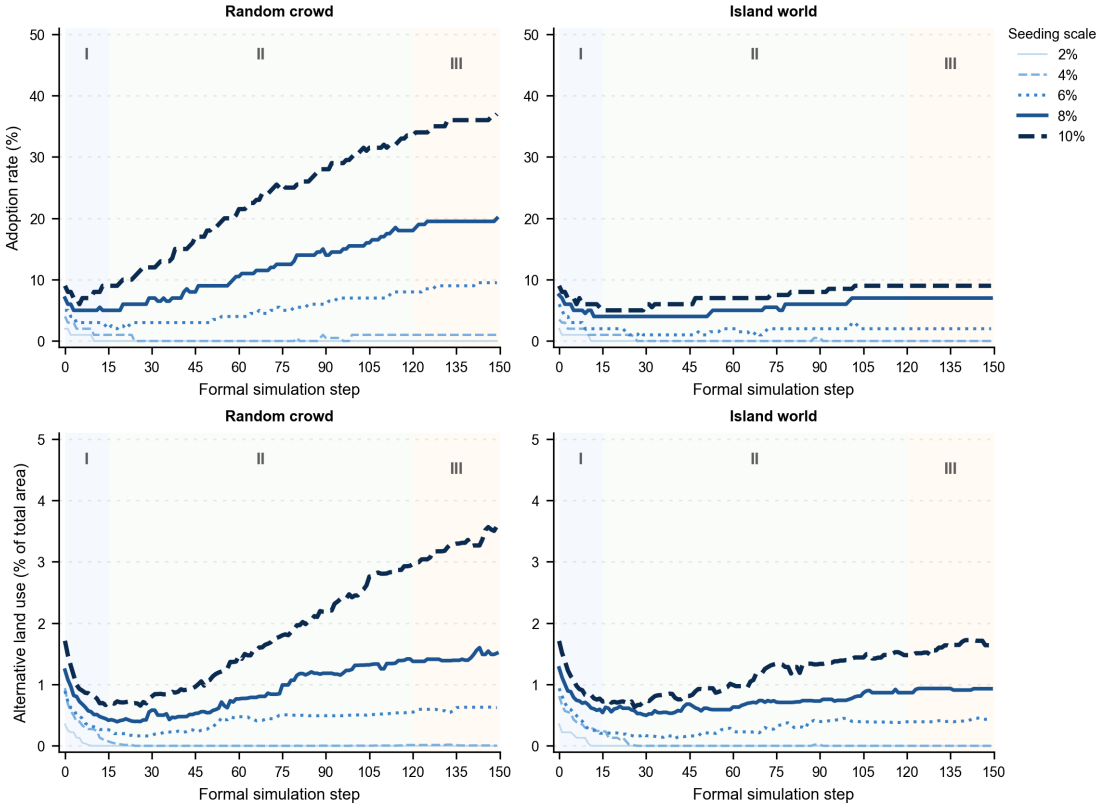

Figure A.1: Adoption dynamics and alternative land use under varying seeding scales across network structures.

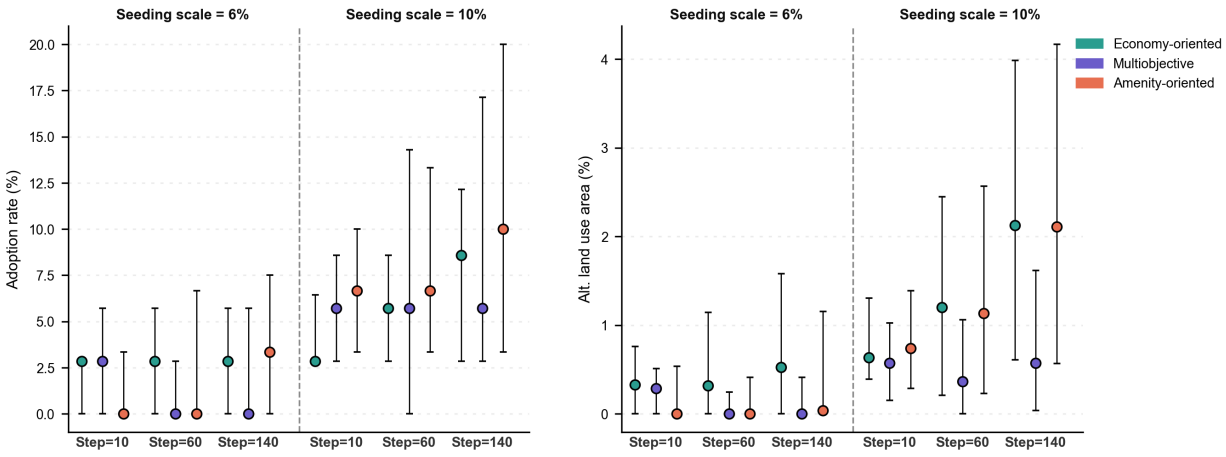

Figure A.2: Adoption and alternative land use by agent type at representative time slices within island world.