TERRA - Trade, Equity Resource Redistribution in Agriculture - An Agent-Based Model for Drought-Prone Areas

Anonymous Author(s)
Submission Id: «submission id»

ABSTRACT

TERRA - Trade, Equity Resource Redistribution in Agriculture – is an agent-based model for agricultural resource management in drought-prone areas. The model simulates decision-making processes of farmers regarding water resource allocation under drought conditions, as well as simultaneously using a game-theoretic approach based on the Ultimatum Game, with "Adaptive Learning Dynamics", such as Reinforcement Learning, Social Learning, and Outcome Adaptation. TERRA's innovative feature is its ability to bridge a significant gap in water resource management (WRM) modelling. Although many studies assess drought impacts or behavioural and game-theoretic responses, few integrate agent heterogeneity, bargaining dynamics, adaptive learning, and policy experimentation within a unified framework. TERRA enables comparative evaluation of interventions, assessing short- and long-term effects on cooperation, agreement stability, and water sustainability in drought-prone agricultural systems. The outputs are used to evaluate the impact of different interventions/policies on sustainability of water resources.

KEYWORDS

Agent-Based Model, Drought-Prone Areas, Water Resource Management, Reinforcement Learning, Social Learning, Outcome Adaptation, Agriculture, Game Theory, Ultimatum Game, Counteroffer Ultimatum Game, Third Party Ultimatum Game

ACM Reference Format:
Anonymous Author(s). 2026. TERRA - Trade, Equity Resource Redistribution in Agriculture - An Agent-Based Model for Drought-Prone Areas. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 8 pages.

1 INTRODUCTION

Water is essential to life and to agriculture, which accounts for roughly 70% of global freshwater withdrawals ([4]). In drought-prone regions, limited and uncertain water availability constrains agricultural productivity and can trigger food insecurity, financial stress, and slower regional development ([7]). Climate change further amplifies these pressures by increasing uncertainty in water availability and intensifying drought risk ([7]). As a result, effective water resource management (WRM) in agriculture is increasingly framed as a problem of both physical scarcity and human decision-making under stress. In shared and scarce water systems, outcomes depend not only on hydrology and infrastructure, but also on how farmers negotiate, cooperate, and respond to perceived fairness in allocation. To analyze such coupled socio-environmental dynamics, this paper adopts a hybrid computational perspective that links behavioural bargaining with decentralized interactions among heterogeneous farmers. Specifically, it introduces TERRA (Trade, Equity & Resource Redistribution in Agriculture), an agent-based model designed to simulate farmers' water-allocation decisions in drought contexts while representing negotiation through a game-theoretic bargaining mechanism inspired by the Ultimatum Game ([6]). A central motivation of TERRA is to help address a persisting gap in WRM modelling: many studies examine drought impacts, or simulate behavioural responses, or apply game-theoretic allocation principles, but fewer integrate (i) agent heterogeneity and local interactions, (ii) bargaining dynamics, (iii) learning and adaptation over repeated interactions (e.g., reinforcement learning, social learning, and outcome-based behavioural adjustment), and (iv) policy experimentation within a single unified decision-support framework. TERRA is intended to support comparative assessment of interventions and policy rules, enabling evaluation of short- and long-term implications for cooperation, agreement stability, and water-resource sustainability in drought-prone agricultural settings.

This paper is structured as follows: In Section 2, we present the literature review, focusing on the main contributions to the field; In Section 3 we introduce Terra, an agent-based model of agricultural water allocation in drought-prone environments. Section 4 contains the implementation of the model, Section 5 addresses the main results and discussion and, finally, in Section 6, we present the main conclusions and limitations of this work.

2 LITERATURE REVIEW

Drought is commonly defined as a prolonged period of abnormally low precipitation that persists for weeks or months and affects large areas ([17]). In agriculture, drought-prone regions are characterized by rainfall uncertainty, degraded soil moisture retention, and limited irrigation access, which together constrain farmers' planning and reduce yields ([3], [9], [13]). These vulnerabilities are increasingly intensified by climate change, as rising temperatures and shifting precipitation patterns contribute to more frequent and severe drought events, with implications for food and water security ([7], [14]). Beyond biophysical exposure, socio-economic conditions, such as poverty and unequal access to irrigation technology, shape the capacity of farming communities to adapt ([4]). Consequently, water resource management (WRM) in drought contexts emphasizes efficiency and resilience through measures such as improved irrigation practices, rainwater harvesting, groundwater recharge, and wastewater reuse, often supported by governance mechanisms that incentivize sustainable adoption and collective management.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus.

Because drought impacts and responses are shaped by heterogeneous actors operating under uncertainty, Agent-Based Models (ABMs) are frequently used to represent decentralized decision-making and emergent system dynamics ([2]). ABMs simulate interactions among farmers, institutions, and environmental conditions, enabling evaluation of policy scenarios and behavioural responses at scale. In shared-resource settings, however, outcomes depend not only on physical scarcity but also on strategic interaction, making Game Theory valuable for formalizing cooperation, conflict, and negotiation under constraints ([15]). A particularly relevant behavioural framework is the Ultimatum Game (UG), a two-player bargaining model in which one agent proposes a division of a resource and the other accepts or rejects the offer, with rejection imposing losses on both sides ([6]). In water allocation dilemmas, UG-inspired mechanisms provide a tractable way to model fairness preferences, bargaining power, and the stability of agreements, and they can be extended through variants such as counteroffers or third-party oversight to better approximate real institutional contexts

To represent adaptation over repeated interactions, "adaptive learning dynamics" can be integrated into ABM-game theoretic systems. Transfer learning concepts motivate this approach by highlighting how knowledge acquired in one environment can inform behaviour in new or evolving conditions ([12]). In practice, three complementary mechanisms are commonly discussed: reinforcement learning (agents learn strategies that improve outcomes over time), social learning (agents imitate successful peers), and outcome adaptation (agents adjust behavioural thresholds in response to recent success or failure). Together, these mechanisms support the simulation of learning and diffusion processes within communities and allow exploration of how negotiation norms and cooperation rates evolve under stress

Recent literature reinforces the relevance of hybrid computational approaches for WRM. Systematic reviews indicate that ABM applications in WRM have expanded substantially since the early 2010s and are increasingly used to model coupled infrastructure and governance challenges across diverse geographic contexts ([1]). In parallel, hybrid studies integrate ABM with cooperative game-theoretic allocation concepts (e.g., Shapley value and least-core solutions) to derive allocations that are both "fair" and stable under hierarchical stakeholder structures ([8]). Advances in multi-agent reinforcement learning propose decentralized coordination frameworks for reservoir networks, emphasizing scalable control rather than explicit bargaining ([5]), while multi-objective reinforcement learning benchmarks frame basin-scale allocation as a trade-off between objectives such as reliability, sustainability, and cost ([10]).

Finally, reviews of ABM in agricultural water trading highlight persistent themes of agent heterogeneity, decision-rule design, uncertainty analysis, and validation—underscoring both methodological maturity and continuing challenges in integrating learning, negotiation, and institutional realism within unified frameworks ([11]).

Synthesizing these strands suggests a continuing opportunity: to connect bargaining-based cooperation mechanisms (capturing fairness and agreement stability) with agent heterogeneity and local interactions (capturing emergent dynamics), while also modelling learning and adaptation across repeated drought-affected seasons and enabling systematic policy experimentation. This motivates the modelling approach developed in the present work.

3 METHODOLOGY

TERRA is as an agent-based model (ABM) of agricultural water allocation in drought-prone environments. To represent land conditions visually and computationally, the environment is discretized into a two-dimensional grid (matrix). Each grid unit encodes the local water reserve state (illustrated in Fig.1 by mapping terrain conditions into discrete reserve levels). This grid-based representation supports both spatial dynamics (e.g., rainfall and drought expansion) and local interactions among neighbouring farmers.

10	7	6	6	7	10	6
7	7	3	7	7	7	7
6	4	2	6	7	6	7
4	3	0	3	6	6	6
6	4	2	4	4	4	4
10	7	6	2	2	2	2
10	7	4	2	0	0	2
7	6	4	3	2	0	2
7	7	4	4	2	2	0
10	7	6	4	3	0	2

10	7	6	6	7	10	6
7	6	3	4	6	7	6
6	4	2	4	7	6	4
4	2	0	3	6	4	4
6	3	0	3	3	3	3
6	4	3	2	2	2	0
7	6	4	2	0	0	0
7	6	4	3	2	0	2
7	7	4	4	2	0	0
10	7	6	4	2	0	2

Figure 1: Before and after: Satellite image shows horrific drought scouring Horn of Africa (Left). Example of a representation of an expansion of the drought-prone area on a matrix (Right)

Farmers are modelled as agents embedded in the landscape, each responsible for a defined portion of the grid. Water allocation between neighbouring farmers is represented using a bargaining mechanism inspired by the Ultimatum Game (Güth et al., 1982). In each interaction, one farmer acts as Proposer (with relatively higher available water) and the other as Responder (experiencing a local deficit). The Proposer offers a quantity of water to transfer, and the Responder either accepts or rejects the offer based on a fairness criterion (implemented as an acceptance rule/threshold). If accepted, the proposed transfer is executed; if rejected, the transfer does not occur and the Proposer retains the transferable water. Repeated interactions across rounds allow the model to represent how bargaining outcomes and cooperation evolve under persistent scarcity.

To capture behavioural change and institutional realism, TERRA extends the baseline bargaining process with additional modules: (i) Ultimatum Game variants, including Counteroffer and Third-Party mechanisms; (ii) Environmental dynamics, including rainfall and drought expansion processes that modify local water reserves over time; (iii) Adaptive learning dynamics, including reinforcement learning (experience-based adjustment of proposer behaviour), social learning (updating strategies by observing successful peers), and outcome adaptation (adjusting behavioural thresholds in response to recent bargaining outcomes); and (iv) policy interventions, such as soil improvement techniques and smart water allocation, which modify environmental constraints and/or decision rules. The model is executed for a predetermined number of rounds, and outputs are used to compare the short- and long-term effects of policies and learning dynamics on negotiation outcomes and water-resource sustainability.

4 IMPLEMENTATION

TERRA was implemented in Python, representing agricultural land as a 36×36 grid (matrix), following a NetLogo-inspired ([16]) visualization approach where spatial patterns and agent interactions can be observed directly. To create meaningful heterogeneity, the grid is partitioned among 144 farmers: each farmer manages a 3×3 block (9 cells), ordered left-to-right and top-to-bottom (F1 at the top-left, F144 at the bottom-right), as shown in Fig.2.

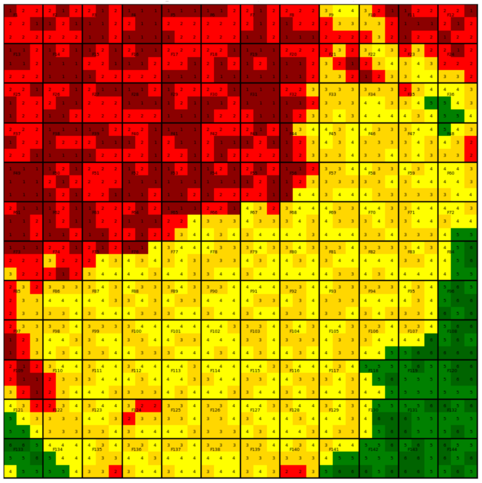

Figure 2: Initial Agricultural Sample Matrix (36 by 36, Seed = 42)

Each cell stores a discrete water-reserve level (the model's "currency" for bargaining and irrigation), ranging from 1 to 6, with an associated colour and land state: 1 (extremely dry, in dark red); 2 (dry, in red); 3 (semi-dry, in dark yellow); 4 (moderate, coloured yellow); 5 (wet, in green) and 6 (extremely wet, in dark green). Higher water levels indicate greater reserves, increasing the likelihood of crop success. All farmers are assumed to grow the same crop type, so water level is the primary determinant of crop viability across plots.

To initialize a drought-prone landscape, the model defines three environmental zones - Dry (1-2), Semi-dry (3-4), and Wet (5-6) - with user-defined target proportions. In the example shown, the target distribution was 30% dry, 40% semi-dry, 30% wet, but the generated matrix produced more dry and semi-dry land and a smaller wet fraction due to spatial constraints (Dry 38.04%, Semi-dry 51.93%, Wet 10.03%).

Terrain generation begins from three user-defined seed coordinates (one per zone) and expands each zone sequentially (Dry → Semi-dry → Wet) using a Breadth-First Search (BFS) filling process. To ensure smooth spatial transitions, the algorithm applies adjacency rules so that neighbouring cells cannot differ too sharply (e.g., a neighbour of level 5 must be chosen from 4,5,6). These constraints promote coherent clusters rather than abrupt jumps (such as 1 directly next to 6), but they also explain why the final zone composition can deviate from the targets. If, after zone targets are reached, some cells remain unassigned, the model fills them using two fallback rules: (i) if neighbouring cells exist, it assigns a compatible value using adjacency constraints; (ii) if no neighbours are defined, it assigns a default random value from the semi-dry range (3 or 4). This guarantees the terrain is fully populated while maintaining a realistic, continuous spatial pattern.

4.1 Ultimatum Game (UG) and variants

TERRA implements a classical Ultimatum Game (UG) to model water-sharing negotiations between farmers. In this setting, both UG roles are farmers: a Proposer offers a transfer of water reserves, and a Responder accepts or rejects the offer. The model's key assumption is that water reserves are transferable between neighbouring farmers. "Water" is operationalized through the discrete water levels stored in each farmer's 3×3 land block. Although random pairing is available (144 farmers → 72 pairs), the model primarily uses spatially local interactions to reflect realistic water-sharing communities. Two neighbourhood definitions are implemented: Von Neumann (N/S/E/W) and Moore (N/S/E/W + diagonals). The experiments reported use the Moore neighbourhood, giving each farmer multiple potential partners (typically 3-8).

For each candidate pair (i, j), the model computes each farmer's total water endowment by summing the water units across their 3×3 block (9 plots). The farmer with the higher total water becomes the Proposer P, and the other becomes the Responder R.

Let W_P and W_R denote the total water of the Proposer and Responder, respectively. The model defines the water gap as

$$\Delta W = W_P - W_R,$$

and the maximum feasible transfer as

$$m = \left\lfloor \frac{\Delta W}{2} \right\rfloor,$$

which ensures that the Proposer cannot become worse off than the Responder after sharing. When ΔW is odd, flooring implies that the Proposer may remain one unit ahead. If $\Delta W < 2$ (equivalently, $m = 0$), the interaction is skipped as non-meaningful (see Fig.3).

Each farmer is characterized by a greed threshold $g \in \{0.1, \ldots, 0.9\}$, which determines the Proposer's offer, and an acceptance threshold $a \in \{0.1, \ldots, 0.9\}$, which determines the Responder's minimum requirement. All offers and requirements are rounded to the nearest integer using half-up rounding, denoted by $\mathrm{round}_{0.5\uparrow}(\cdot)$. The Proposer's offer is defined as

$$\text{Offer} = \max\left\{1,\ \mathrm{round}_{0.5\uparrow}\left((1-g)\,m\right)\right\},$$

which enforces a minimum offer of one unit whenever a bargaining attempt takes place. The respondent's minimum acceptable offer is given by the responder.

$$\text{Requirement} = \mathrm{round}_{0.5\uparrow}\left(a \times m\right).$$

The proposed transfer is accepted if Offer ≥ Requirement, and rejected otherwise, in which case no transfer occurs. When a deal is accepted, the transfer is implemented at the cell level. The Proposer donates water units from the highest-water cells within their 3×3 block (each unit reduces a selected cell by one), while the Responder receives water units into the lowest-water cells of their block (each unit increases a selected cell by one). This mechanism operationalizes an efficient redistribution from relatively water-rich plots to relatively water-poor plots.

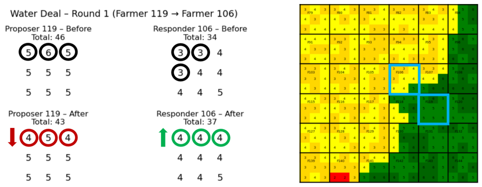

Figure 3: Deal between F119 and F106 (Greed threshold of 0.5 and acceptance threshold of 0.4) + Initial Matrix

To increase institutional realism and diversify bargaining outcomes, two extensions of the ultimatum game (UG) are implemented and can be selected via a simulation parameter.

The Counteroffer UG introduces an additional bargaining step when the Responder rejects the initial offer. Two additional behavioral thresholds are defined. The counteroffer demand threshold d represents the fraction of the maximum feasible transfer m demanded by the Responder when issuing a counteroffer, with higher values of d corresponding to more demanding behavior. The proposer keep threshold k represents the minimum fraction of m that the Proposer insists on retaining when evaluating a counteroffer, with higher values of k indicating stricter behavior by the Proposer.

If the initial offer is rejected, the Responder issues a counteroffer defined as

$$\text{Counteroffer} = \max\left\{1,\ \text{round}_{0.5\uparrow}(d \times m)\right\}.$$

The Proposer evaluates the counteroffer based on the amount they would retain. The Proposer's minimum acceptable retained amount is given by

$$\text{Proposer Requirement} = \text{round}_{0.5\uparrow}(k \times m).$$

The counteroffer is accepted if the retained amount meets or exceeds the Proposer's requirement. If accepted, the counteroffer amount is transferred using the same cell-level transfer mechanism as in the baseline UG. Otherwise, the negotiation fails and no transfer occurs.

The Third-Party UG introduces community oversight into the bargaining process. For each attempted deal, a third farmer is randomly sampled to act as a Judge. The Responder continues to apply the baseline acceptance rule, but the Judge may veto agreements that are deemed unfair.

Two additional elements are introduced. The third-party fairness threshold f specifies the minimum fairness score required for approval. The fairness score is defined as the share of the maximum feasible transfer offered:

$$\text{Fairness} = \frac{\text{Offer}}{m}.$$

The Judge vetoes the agreement if Fairness $< f$.

A deal succeeds only if both conditions are satisfied: the Responder accepts the offer according to their baseline requirement rule, and the Judge approves the agreement by not exercising a veto. This mechanism prevents outcomes that are acceptable to a highly tolerant Responder but are considered exploitative according to community-level fairness standards.

4.2 Environmental Dynamics

Two optional environmental processes capture exogenous variability and stress. A rainfall event deposits a user-defined number of raindrops randomly across the grid. Each raindrop increases the water level of the hit cell by +1, up to the saturation cap (level 6). Drops that hit already-saturated cells are wasted. Rainfall timing (e.g., every r rounds) and intensity (number of raindrops) are configurable. For controlled comparisons across scenarios, rainfall uses a dedicated random seed so that different UG modes can be evaluated under identical rainfall realizations. Drought expansion models a periodic shock that affects all farmers: after a configurable number of rounds, each farmer has one randomly chosen cell within their 3×3 block targeted for drying. If the cell is above the minimum (level 1), it decreases by 1; otherwise, the loss is blocked by the floor constraint. This creates a broad, simultaneous drying pressure that is spatially distributed at the farmer level (see Fig. 4).

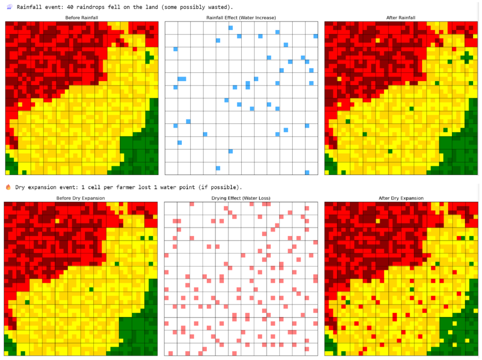
Rainfall event: 40 raindrops fell on the land (some possibly wasted).

Dry expansion event: 1 cell per farmer lost 1 water point (if possible).

Figure 4: Above: Rainfall event (40 raindrops); Below: Dry Expansion (1 cell per Farmer lost 1 water point (if possible))

4.3 Adaptive Learning Dynamics

Adaptive mechanisms are optional and can be activated independently to assess their marginal impact on bargaining success and land outcomes. TERRA implements: Reinforcement Learning (RL), Outcome Adaptation (OA), and Social Learning (SL). RL replaces the greed-based proposer offer rule with a learned offer policy. Each proposer maintains a Q-table over discrete states and actions. Before the main simulation, agents perform a fixed number of one-step bargaining episodes (offline pre-training). Parameters follow standard Q-learning notation: learning rate α=0.1, discount factor $\gamma = 0.9$, exploration rate ϵ=0.2. Because each episode ends after one offer, the update reduces to a terminal one-step form (future term is effectively zero), making Q-values interpretable as smoothed acceptance likelihood estimates for a given (m, t) and action (encode offering 10% to 90% of m).

During the main run, Proposers select offers using a three-tier policy: (i) Own learned state: if the proposer learned (m, t), use its best stored action; (ii) Population-backed fallback: if not learned locally, use the mean best action across farmers for that state, but only if at least a minimum number of farmers learned it (e.g., ≥ 5), to avoid "lucky" outliers; (iii) Default rule: if neither applies, revert to the baseline greed-based offer. This yields a practical balance between individualized learning, robust generalization, and conservative fallback. Outcome Adaptation is a lightweight feedback mechanism: farmers adjust thresholds based on their own recent success rates (no imitation, no value functions). Every fixed interval (e.g., every 2 rounds), thresholds are updated if the farmer has sufficient experience (minimum-event guards prevent overreaction to small samples). Each threshold follows the same structure: compute a rate, then apply small step changes. Typical step sizes used are -0.05 (become less strict after poor outcomes) and +0.02 (become stricter after consistently successful outcomes). OA is applied to several thresholds: Proposer greed (based on proposal failure rate); Responder acceptance (based on responder success rate); Counteroffer demand (counteroffer acceptance as responder); Proposer keep threshold (counteroffer acceptance as proposer); Third-party fairness threshold (judge approval rate). Skipped interactions do not contribute to these statistics, ensuring updates reflect meaningful bargaining exposure.

Social Learning Mechanism

Social learning models the diffusion of successful behaviour across agents. At fixed intervals (e.g., every 5 rounds, after the OA phase), farmers may copy decision thresholds from "role models" who have achieved sufficient success. Eligibility to serve as a role model requires reaching a minimum number of successful outcomes (e.g., ≥ 3).

For each threshold type, an agent updates its value with probability $p = 0.5$. Role models are selected using weighted sampling proportional to their success counts. Specifically successful proposals are used as weights for proposer behaviour; accepted responses are used as weights for responder behaviour; and approved rulings are used as weights for judge behaviour.

To prevent full behavioural homogenization, a small mutation noise term ϵ is added to the copied threshold after updating.

4.4 Policy Recommendations

Two Policy Recommendations were implemented to simulate real decision-making regarding drought prone areas. Soil Improvement Techniques (through soil resistance to drought expansion) represents agronomic practices that increase drought resilience (e.g., mulching, organic matter improvement, terraces, agroforestry, irrigation efficiency). It modifies drought expansion by making losses probabilistic rather than guaranteed. When drought targets a farmer's cell, a Bernoulli resistance test is performed with probability p_{resist}. If the test succeeds, the farmer resists and no reduction is applied. If it fails, the drought loss proceeds as usual (subject to the level-1 floor). The simulation reports interpretable event counts per drought shock: total targeted farmers, resisted cases, realized losses, and floor-blocked cases (already at minimum). In the baseline model, incoming water (after a successful deal) is assigned to the receiver's driest plots first, which reduces extreme dryness broadly. Smart Water Allocation implements an alternative behavioural rule capturing strategic concentration under scarcity: rather than marginally improving many low-productivity plots, a farmer may concentrate water to push some plots into a highly productive state. To avoid unrealistic over-concentration, the implemented strategy is two-phase: allocate water to raise as many plots as possible to level 5 ("Wet"), then allocate remaining water to reach level 6 ("Extremely Wet"). This policy changes where received water is placed without changing how much is received, enabling direct comparisons between distribution strategies at equal transfer totals.

5 RESULTS AND DISCUSSION

5.1 Simulations and Parameters

Five simulations were made: simulation 1 does not include rainfall or policy recommendations. Simulation 2 does not include policy recommendations. The remaining simulations (3-5) are a combination of various policies. All experiments use the same initial landscape to ensure comparability. The matrix is generated with Seed = 42, target zone composition 30% / 40% / 30% (Dry / Semi-Dry / Wet), and seed coordinates (05,05) / (18,18) / (30,30). Due to adjacency constraints, the realized initial composition is 38.04% Dry, 51.93% Semi-Dry, and 10.03% Wet. Across simulations, farmers interact under a Moore neighbourhood over 50 rounds. Baseline bargaining parameters are Proposer greed = 0.8 and Responder acceptance = 0.6. Counteroffer mode additionally uses counteroffer demand = 0.6 and proposer keep = 0.6; Third-Party mode uses a judge fairness threshold = 0.4. Environmental dynamics (when enabled) follow a fixed schedule: drought expansion every 10 rounds (targeting 1 cell per farmer, if possible) and rainfall every 25 rounds with 200 raindrops.
Adaptive learning (when enabled) uses RL pre-training (5,000 episodes; K = 5), Outcome Adaptation every 2 rounds, and Social Learning every 5 rounds. Policy simulations test: Soil Improvement Techniques (SIT) with 30% resistance probability, and Smart Water Allocation (SWA) with a wet target of level 5.

5.2 Introductory Analysis: Standard UG without dynamics (Simulation 1)

Simulation 1 provides a baseline for how the system behaves without environmental stressors, learning, or policies. With high proposer greed (0.8) and a relatively demanding responder threshold (0.6), bargaining success is limited. Most potential interactions are not economically meaningful because neighbouring farmers often have similar water totals. Across 50 rounds: 73.14% of paired interactions are skipped (gap < 2), only 26.86% become attempted deals, and the overall acceptance rate is 21.52% (192 accepted out of 892 attempted deals).
Offers are heavily concentrated at the minimum: most Proposers offer 1 unit, with rare outliers at 2. Cooperation declines quickly: successful deals drop sharply after the first few rounds and then stabilize around 0-1 accepted deals per round, indicating that with fixed thresholds the system reaches a low-cooperation steady regime. Social isolation (unpaired farmers) remains limited but persistent (roughly $6 - 16$ unpaired across rounds). Changes to the landscape are modest. After 50 rounds, Extremely Dry decreases by -1.77 pp (18.13% $\rightarrow$ 16.36%), Wet increases by +2.08 pp (6.33% $\rightarrow$ 8.41%), the Dry zone falls by -0.69 pp, and the Wet zone rises by +0.69 pp. Visually, the main dry/semi-dry/wet clusters remain largely intact. This baseline establishes that standard UG with static behaviour

can redistribute water locally, but only weakly, because most neighbourhood gaps are small and strict thresholds suppress agreement. Without adaptation, the model produces limited cooperation and limited land improvement, even over 50 rounds (see Fig.5).

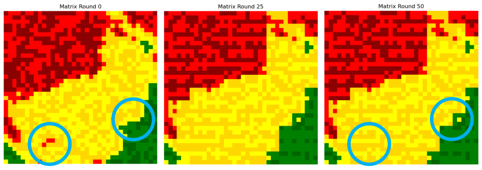

Figure 5: Matrix evolution (initial, 25th and 50th round) – Simulation 1

5.3 UG Type Comparison under environmental stress (Simulation 2, learning OFF)

Simulation 2 introduces drought expansion plus rainfall and compares the three UG types with fixed behaviour (RL/SL/OA OFF). Environmental shocks dominate the baseline trajectory: drought pushes cells downward, rainfall partially compensates, and the net outcome depends on how often bargaining succeeds. In terms of land composition (net change vs initial), in Standard UG, Extremely Dry increases (+0.54 pp), Dry increases strongly (+5.71 pp), while Moderate declines (−5.79 *pp*). Third-Party UG is identical to Standard in this setting (same net composition). Counteroffer UG had the best performance on dryness: Extremely Dry decreases (−3.62 *pp*) and Semi-Dry increases (+5.09 pp) while Moderate declines (−8.18 *pp*).

Acceptance rates clearly separate the modes: Standard: 30.41%, Counteroffer: 59.54%, and Third-Party: 30.41%. Counteroffer nearly doubles the acceptance rate because bargaining adds a second pathway to agreement and, importantly, introduces a "keep" threshold (0.6) that is less restrictive than the baseline greed (0.8). In contrast, Third-Party oversight does not change outcomes here because behaviour is fixed and most failed deals already fail at the responder stage; judge vetoes rarely become the binding constraint. Under drought plus rain, counteroffers materially increase cooperation and reduce extreme dryness, while Standard and Third-Party remain constrained by the initial greed/acceptance mismatch (see Fig. 6).

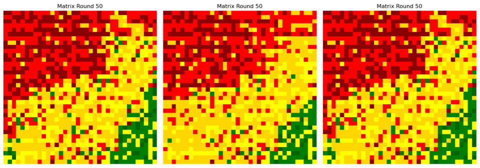

Figure 6: 50th round matrix (Standard, Counteroffer and Third Party, respectively) – Simulation 2 (RL - OFF; SL - OFF; AO - OFF

5.4 Adding Adaptive Learning Dynamics (Simulation 2, learning ON)

Turning RL + SL + OA ON produces the largest shift in system behaviour. The model transitions from low and unstable cooperation to high acceptance rates and sustained redistribution, despite environmental stress.

Cooperation increases sharply: in Standard, acceptance rises from 30.41% to 82.42%; in Counteroffer, 59.54% to 95.83%; and in Third-Party from 30.41% to 73.02%. This is accompanied by fewer skipped deals (gaps increase over time due to shocks and redistribution), and a larger share of attempted interactions (≈ 42-46% attempted deals). In practice, learning makes farmers find workable offers and adjust thresholds so that bargaining does not collapse after early rounds.

Extreme dryness drops substantially. Compared to the initial state (18.13% Extremely Dry), the final Extremely Dry shares become 10.26% (−7.87 pp) in Standard, 8.33% (−9.80 pp) in Counteroffer, and 11.57% (6.56 pp) in Third-Party. The dominant shift is a movement away from the lowest levels into the mid-range (especially Semi-Dry), suggesting that learning-driven bargaining primarily prevents collapse into extreme dryness rather than creating sustained "Wet" regions under repeated drought shocks.

In Standard UG, proposer greed trends downward (less strict), while responder acceptance trends upward (more demanding). The net effect is improved coordination around feasible offers: proposers become more cooperative, and responders raise demands once they observe higher acceptance likelihoods (See Fig. 7).

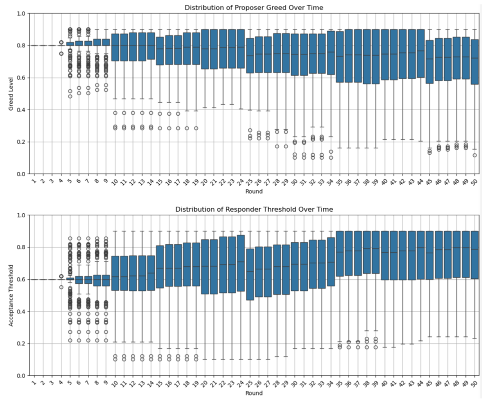

Figure 7: Distribution of Proposer greed (threshold), and Responder acceptance (threshold) over 50 rounds in Standard mode – Simulation 2

In Counteroffer UG, (Fig.8) the strongest driver of improved outcomes is the evolution of the proposer keep threshold, which drops markedly (often toward 0.2), meaning proposers become increasingly willing to accept larger counteroffers even if initial offers remain strict. This "flexibility at the bargaining stage" explains why Counteroffer remains the top performer.

In Third-Party UG (Fig.9), learning also raises the judge's fairness threshold (toward 0.7 − 0.8), making the judge progressively stricter. This institutional constraint reduces acceptance relative

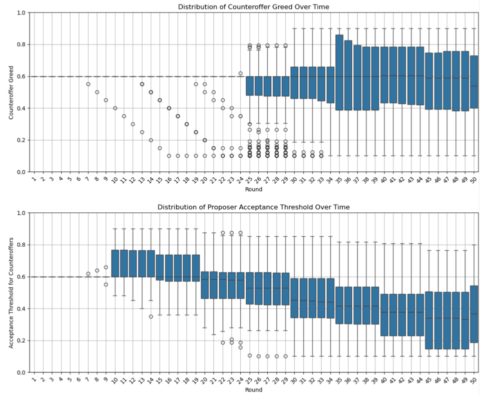

Figure 8: Distribution of Counteroffer demand (threshold), and Proposers keep (threshold) over 50 rounds in Counteroffer mode – Simulation 2

to Standard/Counteroffer and explains why Third-Party is consistently the lowest acceptance mode under learning.

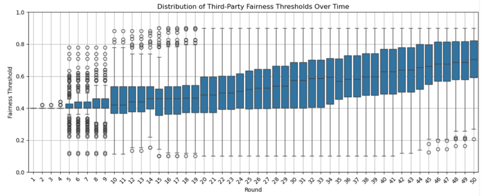

Figure 9: Distribution of third-party fairness threshold over 50 rounds in Third Party mode – Simulation 2

Adaptive learning converts the model from a mostly static, low-cooperation system into a high-cooperation regime, with Counteroffer with learning producing the strongest reduction in Extremely Dry land, and Third-Party producing the most conservative outcomes due to rising institutional strictness (see Fig. 10).

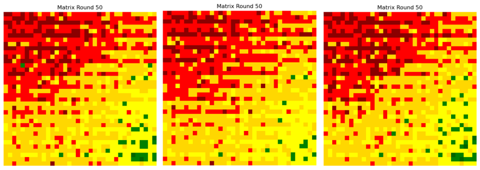

Figure 10: 50th round matrix (Standard, Counteroffer and Third Party, respectively) – Simulation 2 (RL - ON; SL - ON; AO - ON)

5.5 Impact of Policy Recommendations (Simulations 3 to 5)

Policies change the objective from "maximize accepted deals" to "shape land outcomes." Here, SIT targets drought resilience, while SWA targets agricultural productivity by concentrating water into fewer high-quality plots. Across all UG modes, the combined policies produce a clear pattern: Wet land increases strongly (≈ +7.25 to +7.79 pp, reaching ∼ 13.6–14.1%), and Extremely Dry increases relative to the initial level (≈ +1.70 to +2.01 pp, reaching ∼ 20%). This is an important trade-off: SWA reallocates incoming water to push selected plots into wet states (≥ 5), which increases "productive" cells, but it does so by drawing water primarily from Moderate and Semi-Dry areas, leaving a larger share of very low cells under repeated drought. SIT dampens drought losses probabilistically, but in this configuration, it does not fully offset the concentration effect introduced by SWA.

Cooperation remains high but is slightly reduced. Acceptance rates stay strong but fall modestly compared to the learning-only Simulation 2. Acceptance rates are as follows: Standard: 78.00%; Counteroffer: 93.93%; Third-Party: 71.53%. This suggests that policies primarily reshape where water goes after transfers, rather than fundamentally improving bargaining feasibility.

A scan of all 96 combinations shows two "best-case" directions depending on the objective:

- **Minimizing the Dry zone:** The top results come from Simulation 3 (SIT) with RL + SL enabled (OA optional), achieving Dry zone = 35.34% (−2.70 pp). However, these configurations do not preserve Wet zone share, indicating that SIT primarily prevents deterioration rather than building wet clusters.
- **Maximizing the Wet zone:** The top results come from Simulation 5 (SIT + SWA), achieving Wet zone ≈ 17.90% (+7.87 pp), but with a marked increase in Dry zone ≈ 42.67% (+4.63 pp). This confirms that boosting high-productivity plots can come at the cost of broader dryness.

Policies introduce a controllable trade-off between (i) reducing dryness (resilience-oriented) and (ii) creating productive wet plots (concentration-oriented). In this model, SIT supports dryness reduction, while SWA supports wet-plot creation, and combining them prioritizes productivity gains even if extreme dryness increases (see Fig. 11).

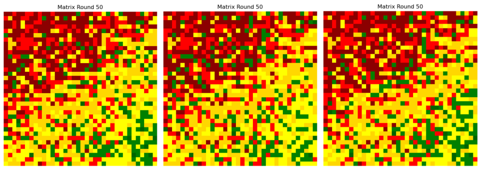

Figure 11: 50th round matrix (Standard, Counteroffer and Third Party, respectively) – Simulation 5 (RL - ON; SL - ON; AO - ON, SIT - ON, SWA - ON)

6 CONCLUSIONS AND LIMITATIONS

Across the staged experiments, results show a clear progression. With fixed behaviour, the standard UG yields limited redistribution, largely because many neighbour pairs have small water gaps and strict thresholds suppress agreement.

Introducing counteroffers substantially increases acceptance under environmental stress, demonstrating how bargaining flexibility can sustain cooperation during drought. When adaptive learning is enabled, the system shifts to a high-cooperation regime and extreme dryness declines markedly, indicating that endogenous behavioural adjustment can be as influential as the choice of UG variant. Finally, policy modules highlight an explicit objective trade-off: Soil Improvement Techniques (SIT) support drought resilience (reducing exposure to drying shocks), whereas Smart Water Allocation (SWA) concentrates incoming water to increase the share of productive wet plots—sometimes at the expense of broader dryness. Together, these mechanisms position TERRA as a decision-support environment for testing how institutions, learning, shocks, and policies interact in drought-prone agricultural systems.

Overall, TERRA demonstrates that an agent-based representation of farmers can be meaningfully coupled with game-theoretic bargaining to study how fairness constraints, negotiation rules, and enforcement mechanisms shape water redistribution under scarcity. Environmental dynamics provide exogenous stressors, while policy levers allow controlled intervention tests. The combination of standard UG, counteroffer, and third-party variants supports comparisons between baseline bargaining, negotiated compromise, and community oversight, enabling more nuanced interpretation of cooperation, exploitation risk, and institutional strictness.

Several limitations motivate future work. First, reinforcement learning currently targets proposer offers via offline pre-training: learning is therefore asymmetric across roles and not fully embedded throughout the full simulation horizon. Extending learning to responder behaviour (e.g., acceptance thresholds), and/or adopting online and multi-step learning for bargaining sequences, would strengthen behavioural realism. Second, because transfers are discrete and constrained by small water gaps, many distinct greed settings collapse to the same minimum integer offer, reducing behavioural resolution and limiting the model's ability to differentiate strategies in near-equal neighbour conditions. This can be mitigated by introducing finer-grained resource units, alternative transfer caps, or continuous/lot-size offers, which would better preserve strategic variation in low-gap settings.

REFERENCES

[1] K. Aybuğa and A. G. Yücel Işıldar. 2022. Agent-Based Approach on Water Resources Management: A Modified Systematic Review. *Turkish Journal of Water Science and Management* 6, 2 (2022), 202–236. https://doi.org/10.31807/tjwsm.1123808

[2] Thomas Berger and Christian Troost. 2014. Agent-based modeling of climate adaptation and mitigation options in agriculture. *Journal of Agricultural Economics* 65, 2 (2014), 323–348. https://doi.org/10.1111/1477-9552.12045

[3] Aiguo Dai. 2011. Drought under global warming: A review. *Wiley Interdisciplinary Reviews: Climate Change* 2, 1 (2011), 45–65.

[4] Food and Agriculture Organization of the United Nations. 2017. *The Future of Food and Agriculture: Trends and Challenges.* FAO, Rome.

[5] H. Fu, G. Xiong, J. Li, and S. Lin. 2025. Multi-agent reinforcement learning for decentralized reservoir management via murmuration intelligence. arXiv preprint arXiv:2504.11569. https://arxiv.org/abs/2504.11569

[6] Werner Güth, Rolf Schmittberger, and Bernd Schwarze. 1982. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization* 3, 4 (1982), 367–388.

[7] Intergovernmental Panel on Climate Change. 2021. *Climate Change 2021: The Physical Science Basis.* Cambridge University Press. https://doi.org/10.1017/9781009157896

[8] M. S. Khorshidi, M. R. Nikoo, G. Al-Rawas, N. Bahrami, M. Al-Wardy, N. Talebbeydokhti, and A. H. Gandomi. 2024. Integrating agent-based modeling and game theory for optimal water resource allocation within complex hierarchical systems. *Journal of Cleaner Production* 482 (2024), 144164. https://doi.org/10.1016/j.jclepro.2024.144164

[9] Ashok K. Mishra and Vijay P. Singh. 2010. A review of drought concepts. *Journal of Hydrology* 391, 1-2 (2010), 202–216.

[10] Z. Osika, R. Rădulescu, J. Zatarain Salazar, F. Oliehoek, and P. K. Murukannaiah. 2025. Multi-objective reinforcement learning for water management. arXiv preprint arXiv:2505.01094. https://arxiv.org/abs/2505.01094

[11] S. Ozkal, E. Bertone, and R. A. Stewart. 2025. A systematic review of agent-based modelling in agricultural water trading. *Water* 17, 6 (2025), 869. https://doi.org/10.3390/w17060869

[12] Sinno Jialin Pan and Qiang Yang. 2010. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering* 22, 10 (2010), 1345–1359. https://doi.org/10.1109/TKDE.2009.191

[13] Johan Rockström, Malin Falkenmark, Louise Karlberg, Holger Hoff, Stefan Rost, and Dieter Gerten. 2009. Future water availability for global food production: The potential of green water for increasing resilience to global change. *Water Resources Research* 45, 7 (2009), W00A12. https://doi.org/10.1029/2007WR006767

[14] Kevin E. Trenberth, Aiguo Dai, Gerard van der Schrier, Philip D. Jones, Jonathan Barichivich, Keith R. Briffa, and Justin Sheffield. 2014. Global warming and changes in drought. *Nature Climate Change* 4, 1 (2014), 17–22.

[15] John von Neumann and Oskar Morgenstern. 1944. *Theory of Games and Economic Behavior.* Princeton University Press, Princeton.

[16] Uri Wilensky. 1999. NetLogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL. http://ccl.northwestern.edu/netlogo/

[17] World Meteorological Organization. 2021. *2021 State of Climate Services: Water.* Technical Report WMO-No. 1278. World Meteorological Organization, Geneva.