

Individual Decision-Making in Coupled Agent-Environment Systems

Karolina Chłopicka (15476758) & Victoria Peterson (15716546) &
Shania Sinha (14379031) & Henry Zwart (15393879)

Abstract—The stringent climate targets being set on global and local levels underscore a need to understand how individuals make decisions that determines environmental outcomes. Traditional mean-field and two-strategy frameworks often overlook individual-level cognitive processes, dynamic risk perception, and evolving social norms that drive real-world behavioral change. To address these gaps, we develop a spatially explicit agent-based model (ABM) in which heterogeneous agents choose between climate-friendly and degradative actions based on adaptive utility functions incorporating intrinsic environmental preference, local social pressures, and perceivable ecosystem feedback. We integrate novel sub-models for dynamic risk-perception thresholds, shifting-baseline effects, and memory-based spillover. We also perform a mean-field analysis alongside a PAWN sensitivity analysis. Phase plots and clustering analyses reveal critical tipping points and emergent patterns — ranging from multi-stability and oscillatory regimes to scale-free cluster formation — under varying policy levers such as varying rationality and rate of adaptation to the environment. Our results quantify the threshold levels of key parameters required to shift a socio-ecological system from low-cooperation, degraded equilibria toward sustainable, high-cooperation trajectories, providing actionable insights for designing robust, climate-positive interventions.

I. INTRODUCTION

In recent years, global climate policy has coalesced around ambitious national targets for greenhouse-gas reductions. Recently, the European Union’s Fit for 55 legislative package legally mandated a reduction in net GHG emissions of at least 55 percent by 2030 relative to 1990 levels, paving the way for climate neutrality by 2050 [1]. With other major emitters setting stringent targets to achieve a positive climate outcome, it is increasingly important to understand and explore the complexity and dynamism underlying human behaviour in social-ecological systems. Insights into how such policy targets translate into individual and collective environmental outcomes can be achieved by modeling the coupled-dynamics of human decision-making and ecosystem processes using Agent-Based Models (ABMs).

Although Kraan et al. (2019) [2] developed ACT — an ABM that departs from mean-field assumptions — to explore how heterogeneity, leadership, and local network structures drive critical transitions in low-carbon energy systems, it does not account for dynamic environmental feedback on individual risk perception or the evolving influence of social norms. Conversely, Tilman et al. (2020) [3] formulated a general eco-evolutionary game framework that captures how strategic choices and environmental states mutually shape one another. However, its mean-field, two-strategy abstraction overlooks

agent-level cognitive processes such as shifting baselines, pluralistic ignorance, and memory-based spillovers. It also does not consider network heterogeneity, that is crucial to real-world decision-making under climate change. Furthermore, previous critiques of ecological ABMs call for the integration of formal decision-making theories to enrich the psychological foundation of these models [4]–[6]. Thus, a key research gap remains in integrating these nuanced psychological factors into eco-evolutionary models to better represent how diverse individual decision rules, social influences, and dynamic risk perceptions jointly determine collective environmental outcomes.

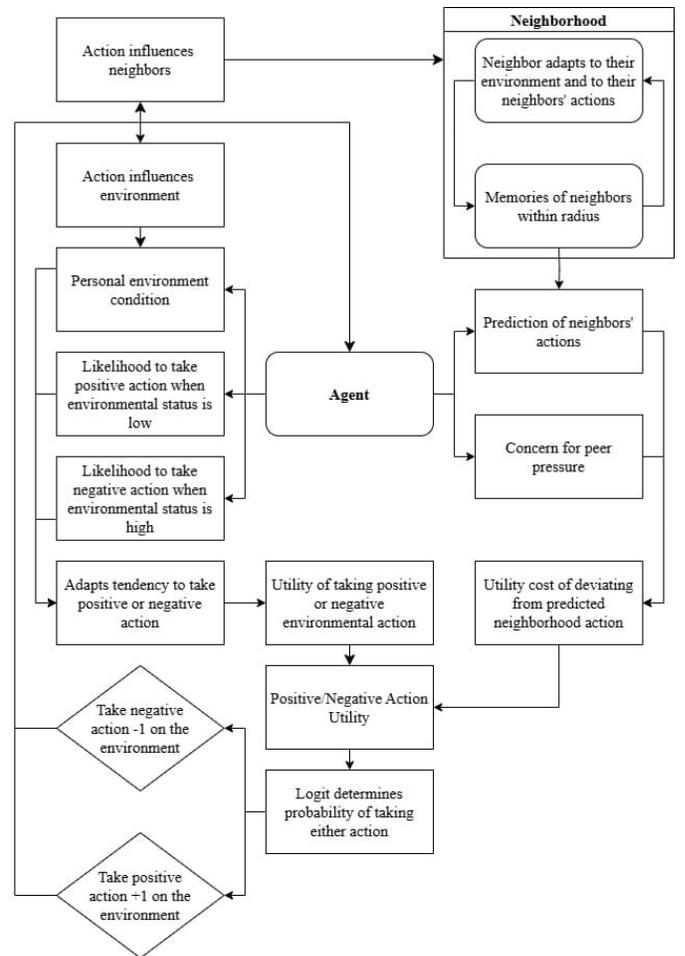


Fig. 1: Schematic representation of the interactions between an agent, its neighbors and its environment in a single timestep.

In this paper, we aim to investigate the following research questions:

- 1) *How do the steady-state equilibria differ in the coupled socio-ecological model?*
- 2) *What role do social-influence structures play in emergent cooperation?*
- 3) *What dynamical regimes appear in the ABM but are absent (or smoothed-out) in its mean-field approximation?*

We begin in Section II by surveying foundational work on social norms, coupled human–environment systems, and agent-based climate models. Section III introduces our model and its mean-field analysis, complete with utility functions and dynamical equations. Section IV outlines our PAWN sensitivity analysis, identifying the parameters that most shape outcomes. Section V presents phase plots and critical transition analyses, revealing tipping points and cluster formation. It also discusses the implications of our findings, and their relation to the mean-field model. Finally, Section VI concludes with directions for future exploration.

II. RELATED WORK

A. Climate-related behavioural factors

Despite the ever-increasing severity of climate change, both individual and population-level environmental behaviour remain complex and difficult to predict. Existing research into the factors influencing climate-related decision-making is highly interdisciplinary, with theoretical and empirical studies spanning environmental science, psychology, economics, tourism, agriculture, mathematics, and computational science. Central to many theoretical studies is the Theory of Planned Behaviour [7]–[9], a general theory describing human decision-making as a combination of individual preferences, social norms, and ability/readiness to take a given action.

The impact of social norms on climate-related behaviours has been widely demonstrated theoretically [10]–[12], empirically [9], [12], [13], and computationally [8], [14]. A double-edged sword, social norms can serve both to deter and to sustain climate-friendly behaviour. Historical inaction on climate mitigation can limit the adoption of climate-friendly behaviours even when this is favoured by current individual preferences — a phenomenon known generally as *pluralistic ignorance* [12], [15], [16].

Individual perceptions of climate-change severity and risk are shown to reinforce willingness to adopt climate-friendly behaviour when environmental degradation is readily visible (e.g., severe weather events, or agricultural impacts) [17]–[22]. A stated willingness to adopt, however, does not necessarily translate to realised behavioural change [23]–[25]. Uncertainty, risk perception, and trust in science are demonstrated to influence decision-making [26], [27], however, their effects can be nullified by the prevalence of opposing social norms. Individual optimism regarding the effectiveness of climate action has also been shown to slow its adoption [28]. Computational studies have demonstrated an increased likelihood of behaviour consensus resulting from local interactions [29].

The influence of perceived climate risk on individual preferences for climate-friendly behaviour is not necessarily linear.

Several studies [10], [30] argue that individuals are most likely to adopt climate-friendly behaviours when environmental degradation is visible, but *only when* it is also perceived as reversible. Climate-friendly behaviour is de-prioritised when the environment is perceived as healthy [29], [31], and its efficacy is questioned when the environment is sufficiently degraded [18], [32]. Additionally, the *shifting baseline effect* suggests a gradual decline in perceived climate change severity — even when the true state continues to worsen — due to the establishment of a new status quo [33]–[36]. External factors such as inequality further limit adoption [37], [38].

Finally, present-day behavioural preferences of individuals can directly influence their future behaviour. Examples include the *green spillover* and *green licensing* effects, which suggest that current climate-friendly actions can positively or negatively reinforce future behaviour, with individual perception theorised as a mediator [13], [39]–[44].

B. Coupled human-environment systems

As discussed in the previous section, perception of environmental degradation is broadly-accepted as contributing to individuals’ decision-making regarding behaviours relating to climate change. Recent years have seen an acceleration of research into *Coupled Human-Environment Systems* (CHES), studying the feedback loop that arises when one considers how agents’ actions in turn affect the environment [45]–[49]. Building on prior work by Weitz et al. (2016), [50], Tilman et al., (2020) [3] proposed an evolutionary game theoretic framework for studying common resource dilemmas with dynamic environments. This work has been extended to examine interplay with different agents behaviours [51], implications for sustained climate action, and risks such as tipping points [30], [49].

C. Agent-based climate behaviour modelling

Agent-based research into climate-related behaviour is a rich and expanding landscape, and as such we focus here on work identified as most relevant to the research question outlined in I. For a more complete overview, we invite the reader to consult one of several recent surveys of the wider field [8], [52], [53].

The modelling framework introduced by Tilman [3] is not inherently agent-based, but has served as a launchpad for ABM studies with heterogeneous attributes [51]. ABM studies of coupled human-environment system ABM studies also feature opinion dynamics [54], [55], as well as the effects of delayed feedback on equilibria stability [56]. Numerous studies examine the effects of social norms on climate-related behaviour, including investigating energy consumption [57], pluralistic ignorance [58], evolving agent preferences [59], or critical transitions in public support for climate mitigation [2], [60], [61].

The interdisciplinary insights, expounded in the previous sections, provide a basis for an ABM that operationalizes the Theory of Planned Behaviour by balancing intrinsic environmental preferences, CHES-inspired social-norm pressures

(including pluralistic ignorance), and perceived action feasibility. Dynamic risk-perception thresholds produce peaks in eco-actions only when environmental harm is both visible and reversible, while shifting-baseline effects progressively diminish engagement. Memory-based spillover and licensing mechanisms allow past behaviours to reinforce or undermine future cooperation, and coupled human–environment feedback loops tie collective actions back to changes in the environmental stock. Incorporating heterogeneity in social influence, cognitive parameters (memory length, forecasting ability, rationality), and network topology yields emergent patterns such as cooperator clustering, oscillatory engagement cycles, multi-stable equilibria, and ecological tipping points between collapse and recovery.

III. THEORY

We developed an agent-based model for human decision-making in a dynamic environment, in which individuals’ behavioural choices are both *informed by* the environment and *affect* the environment. In this section we outline fundamental model mechanics, and summarise key analytical results from our corresponding mean-field analysis. A complete model description is provided in ODD-D format in Appendix A, followed by the mean-field derivations in Appendix B.

A. Environmentally-influenced decision-making

Consider a population of $k \in \mathbb{N}$ individuals, $N = \{1, \dots, k\}$, who repeatedly choose between adopting **Climate-friendly** (C) or **Degradative** (D) behaviours. An individual i ’s decisions depend on their own perceptions of both the environment (n_i^t) and the social norms imposed by their neighborhood of influence (\bar{a}_i^t).¹ The former can amplify an individual’s preference for climate-friendly behaviour when the environment is viewed as *damaged but reversible*, but also dampen it when the environment is particularly healthy [10], [30]. Social norms provide stability in decision-making via the status quo, with potential to both slow and sustain climate-friendly behaviour [9]–[13]. We model an agent i ’s decisions using a discrete-choice logit model, with homogeneous rationality parameter λ , and representative utility V_i ,

$$\begin{aligned} V_i(C) &= (1 - w_i) \cdot s_i - w_i \cdot (1 - \bar{a}_i)^2 \\ V_i(D) &= (1 - w_i) \cdot (4 - s_i) - w_i \cdot (1 + \bar{a}_i)^2 \end{aligned} \quad (1)$$

Where $s_i \in [0, 4]$ is i ’s level of preference for climate-friendly behaviour and $w_i \in [0, 1]$ is a static attribute which moderates the relative contributions of individual preference and social pressures to i ’s decisions. The probability that i chooses an action $a \in \{D, C\}$ is given by:

$$\mathbb{P}(a_i^t = a) = \frac{\exp(\lambda V_i(a))}{\exp(\lambda V_i(C)) + \exp(\lambda V_i(D))} \quad (2)$$

¹We henceforth omit the t superscript for visual clarity, except where doing so would cause ambiguity.

B. Coupled human–environment dynamics

Agents’ action preferences and local environment states vary with respect to one another according to the dynamical system in Equations 3 and 4 respectively. The function σ is the logistic map, $\sigma : n \mapsto 4n(1 - n)$. The relative rates at which the environment is restored/degraded, and preference for cooperation/defection increase are set by parameters β_n^+, β_n^- and β_s^+, β_s^- , with the overall environment and preference update speeds defined by γ_n and γ_s .

$$\frac{1}{\gamma_s} \frac{ds_i}{dt} = \beta_s^+ \cdot \sigma(n_i^t)(4 - s_i^t) - \beta_s^- \cdot (1 - \sigma(n_i^t))s_i^t \quad (3)$$

$$\frac{1}{\gamma_n} \frac{dn_i}{dt} = \beta_n^+ \cdot (1 - n_i^t)\alpha_i^t - \beta_n^- \cdot n_i^t(1 - \alpha_i^t) \quad (4)$$

All experiments in this report (with the exception of the sensitivity analysis) take $\beta_n^+ = \beta_n^- = \beta_s^+ = \beta_s^- = 1$

C. Mean-field model

We derive expressions for the expected mean action its temporal dynamics via a mean-field analysis, making the following simplifying assumptions:

- 1) $w_i = \mathbb{E}[w_i] := w$ for each agent $i \in N$,
- 2) $\frac{dP}{dt} = 0$ for each $i \in N$, such that $s_i^t = s$, and
- 3) $\mathbb{P}(a_i, a_j) = \mathbb{P}(a_i) \cdot \mathbb{P}(a_j)$ for each pair of agents $i, j \in N$.

Expanding subterms in Equation 2 yields a simplified expression for an agent’s expected action:

$$\mathbb{E}[a_i] = \tanh(\lambda[(1 - w_i)(s_i - 2) + 2w_i\bar{a}_i]) \quad (5)$$

We extend this to derive a fixed-point expression for the population-level expected action, $m := \mathbb{E}[a]$,

$$m = \tanh(\lambda \overbrace{[(1 - w)(s - 2) + 2wm]}^z) \quad (6)$$

which has between 1 and 3 solutions, at most two of which are stable. The three-solution case corresponds to pluralistic ignorance, where the influence of social norms is sufficient to sustain a behavioural state different from the mean agent preference. For this to occur, it is necessary (though not sufficient) that $\frac{1}{2\lambda w} < 1$ — that is, the expected social pressure or rationality must be sufficiently large. We obtain the dynamics of m as the time derivative of Equation 6:

$$\frac{dm}{dt} = \frac{(1 - w)\lambda \operatorname{sech}^2(\lambda z^t)}{1 - 2\lambda w \operatorname{sech}^2(\lambda z^t)} \cdot \frac{ds}{dt} \quad (7)$$

IV. SENSITIVITY ANALYSIS

When the model has multiple parameters, the sensitivity analysis is key to distinguishing influential parameters from the ones that don’t significantly impact the model’s behavior. This, in turn, can be used to simplify the model by removing such parameters. Additionally, sensitivity analysis can help identify critical or otherwise interesting regions in the space of the input factors. A suitable sensitivity analysis must be chosen with careful attention to the characteristics of the model output. Failing to do so can undermine the reliability of the results.

As model input, we chose to investigate the following parameters: Grid length, rationality (λ), memory size, adaptation rate (γ_s), neighborhood radius, and environment recovery rate (β_n^+). The range of values used for parameter sampling is presented in Table I.

Parameter	Min. value	Max. value	Type
Grid length	5	50	integer
Rationality	0	10	float
Memory size	2	10	integer
Adaptation rate	0.001	0.05	float
Neighborhood radius	0	1	integer
Recovery rate	0.5	2	float

TABLE I: Range of parameter values used in the PAWN sensitivity analysis.

Following the procedure for global sensitivity analysis parameter space was sampled, and model simulation was performed in order to gather output statistics. Samples were generated using Saltelli’s extension of the Sobol sequence (as implemented in the SALib Python package). For our simulation, the number of samples was set to $N = 2048$ and the number of parameters $D = 6$. This results in $N(D + 2)$ simulations [62]. Models were re-run 5 times for each parameter set, with different random seeds to account for stochasticity, for a total of 81920 simulations:

$$(2048 \cdot (6 + 2)) \cdot 5 = 81920$$

We chose 6 measures of the model output:

- *Mean environment* - average state of the environment at the final time step, computed over all cells of the 2-D lattice.
- *Mean action* - average agent’s action at the final time step, computed over all cells of the 2-D lattice.
- *Pluralistic ignorance* - capturing how much an agent’s behavior is distorted due to the misperception of social norms.
- *Cluster count* - number of healthy environment clusters at the final time step. Cell was considered healthy if the mean value of the environment was above 0.6.
- *Peak frequency* - the frequency at which oscillations have the highest average power within its full-width.
- *Dominant frequency power* - measure of the oscillation mode with the largest magnitude or impact on system stability.

To ensure measurements are taken at system equilibrium, we run each model for 2000 timesteps, which our testing shows to be sufficient in most cases.

The empirical distributions of the output metrics are shown in Figure 2. These types of distributions are not suitable for variance-based sensitivity methods because such methods rely on variance being a meaningful measure of output uncertainty. However, this assumption does not hold for multi-modal or highly skewed distributions. For this reason, we chose to use the PAWN method for our sensitivity analysis. Unlike variance-based approaches, PAWN is a density-based method that can be applied effectively to all types of output distributions, including those that are highly skewed or multi-modal [63].

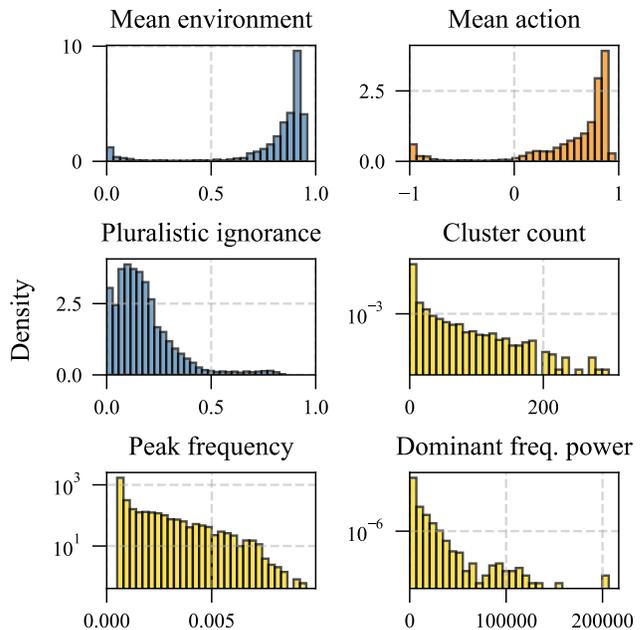


Fig. 2: Empirical output distributions as observed across all samples.

The results of the PAWN sensitivity analysis are shown in Figure 3. Each row corresponds to an output metric of the model, while each column represents one of the input parameters. The values in the heatmap indicate the importance of each parameter for the corresponding output.

Several clear patterns emerge. Mean environment and peak frequency are most influenced by memory size. Mean action is primarily affected by recovery rate, and the same is true for Pluralistic ignorance. Cluster count and dominant frequency power are most influenced by adaptation rate.

Neighborhood radius has minimal influence on any of the selected output metrics, indicating that vision structure has limited impact on the system’s overall behavior. This may be caused by the fact that only two values of radius are available in the model, one for the Moore neighborhood and one for the full grid. Similarly, grid length shows low sensitivity across most outputs, making it the second-least impactful parameter. Overall, this suggests that the spatial structure of the grid plays a relatively minor role compared to other parameters.

V. RESULTS & DISCUSSION

A. Phase Plots

By plotting the continuous values of rationality λ and the rate of change of action preference γ_s on heatmaps of different outputs, we are able to observe general behavioral transition points that may act as critical phase transitions of the system.

The model setup specified in Figure 4 for the experiment without dynamic action preference values nor neighborhood action prediction in Figure 4a shows how the base model expresses oscillatory behavior through Fourier analysis of the dominant oscillation power. No critical transition exists in this behavior, but we generally see how as rationality approaches

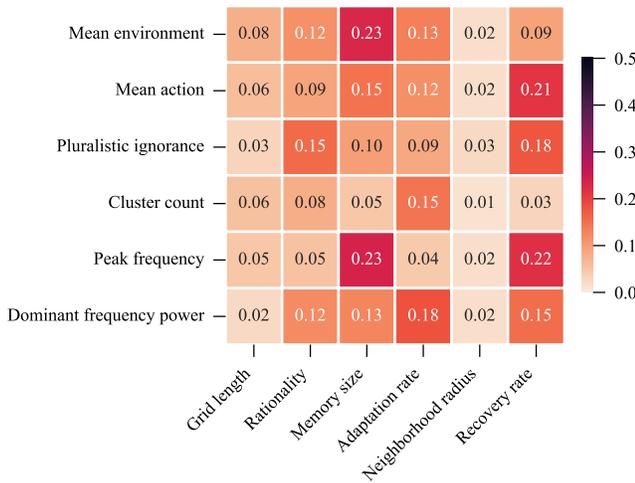
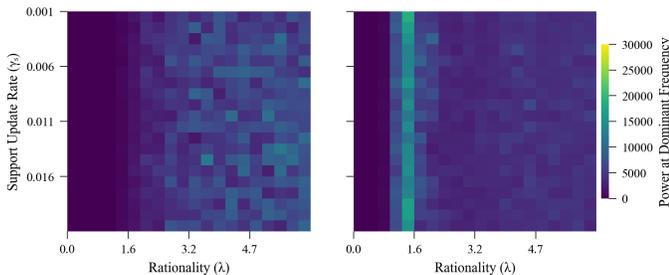


Fig. 3: First-order PAWN sensitivity indices.

0, rational model dynamics become overshadowed by random discrete choices, reflecting in the Fourier analysis as a lack of any oscillation and in the environmental status as random noise centered around 0.5. As the rationality increases, we observe how the natural feedback of environment status and probabilistic agent actions creates noisy periodic behavior.



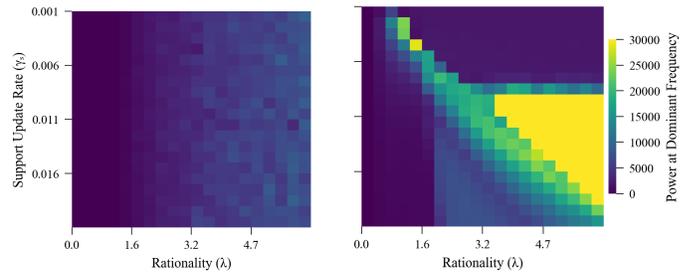
(a) No neighborhood predictor (b) Linear neighborhood predictor

Fig. 4: Fourier oscillation strength analysis for non-dynamic action preference, single radius neighborhood, and 1.0 peer pressure coefficient for all agents

Modifying this same setup to include a linear regression predictor sub-model over neighbors' past 10 actions in Figure 4b shows how a predictive model adds 'momentum' to neighbors' action influence by smoothing short-term fluctuations and reinforcing established behavioral trends, reducing the sudden social pattern shifts that exist in the non-predictive averaging algorithm. This new tendency causes oscillations to be dampened as rationality increases with a notable exception around $\lambda \approx 1.25$.

Further exploration shows the cause of the spike by looking at the parameters used in Figure 5a, a setup which uses a peer pressure coefficient initialization of random values between 0.0 and 1.0 to create a smooth transition across rationality instead of the uniform 1.0 across all agents. This may mean that the "momentum" expressed by the neighborhood prediction sub-model creates a sudden spike in environmental status

when all agents place equal weight on social pressures, but a randomized initialization reduces this by adding noise to the agent parameter.



(a) Linear neighborhood predictor, no dynamic action preference (b) Dynamic action preference, no neighborhood predictor sub-model

Fig. 5: Fourier oscillation strength analysis for single radius neighborhood, and random uniform peer pressure coefficient between $[0.0, 1.0]$ for all agents

In these previous cases, we've observed oscillations of system-wide behavior, but no emergent local behaviors or critical phase transitions. The introduction of the dynamic action preference sub-model reveals complex, parameter-dependent behaviors, as shown in Figure 5b.

Qualitatively, the adaptive model has three distinct sections:

- A diagonal spike in oscillation strength
- A "square" of oscillation activity in the range greater than $\gamma_s \approx 0.08$ and $\lambda \approx 2.0$
- A distinct separation of low and high oscillation within the "square", divided by the diagonal.

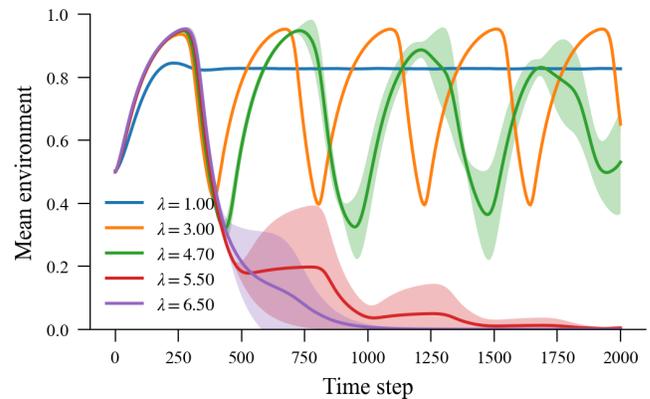
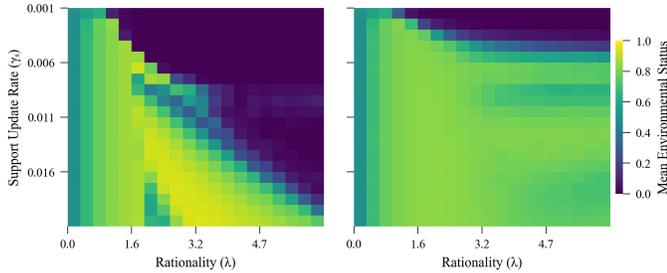


Fig. 6: Environmental status time series of $\gamma_s = 0.01$ with dynamic action preference, no neighbor prediction, single radius neighborhood, and random uniform peer pressure coefficient between $[0.0, 1.0]$ for all agents

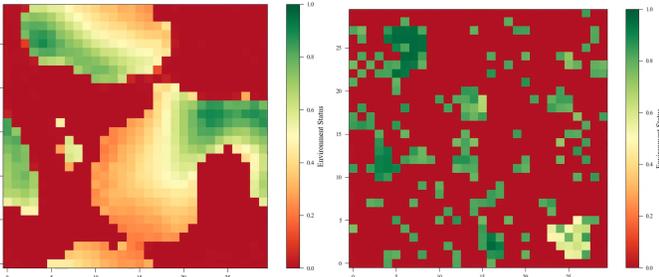
By selecting a specific value of $\gamma_s = 0.01$ and taking the average environmental status output at 2000 time steps in Figure 6, we begin to see the possible distinctions between each section of the heatmap. The darkened portion of the "square" with $(\lambda = 3.0)$ represents seemingly stable oscillations with very little variance, while the diagonal portion

($\lambda = 4.7$) is an unstable oscillation with large variance. The bright yellow section indicates a high oscillation strength, which is caused by the singular large spike that collapses into an equilibrium ($\lambda = 5.5, 6.5$), while the surrounding areas of 0 oscillation strength are areas where no significant oscillation occurs before reaching a non-zero equilibrium ($\lambda = 1.0$).

This is further clarified by Figure 7a, where the final environmental state at time 1000 reaches a positive equilibrium to the left of the "square", a zero equilibrium above the diagonal, a cyclical equilibrium below the diagonal, and noisy oscillations on the diagonal.



(a) 1.0 social pressure weight for all agents (b) Random social pressure weight for all agents



(c) Environment cluster formation with 1.0 social pressure weight for all agents, $\gamma_s = 0.009$ and $\lambda = 3.2$ (d) Environment cluster formation with random social pressure weight for all agents, $\gamma_s = 0.004$ and $\lambda = 6.0$

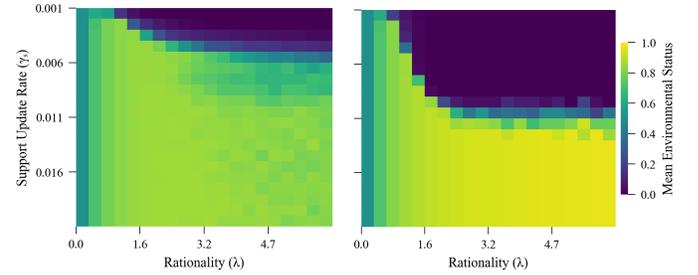
Fig. 7: Environmental status at time step 1000 for an dynamic action preference, no neighborhood action prediction, and single radius neighborhood

If we visualize these noisy oscillations along the diagonal in Figure 7c, we find that the system begins to express some emergent clustering behavior, where portions of the population stabilize on negative environmental status and negative action, while the others continuously oscillate their decision-making erratically.

In the case where agents' place non-uniform weight on social influence, Figure 7b expresses a similar influence as was observed in Figure 5a, where-in the original uniform influence of 1.0 created large spikes in oscillatory behavior, which are no longer present in the new model. Indeed, a heatmap of dominant oscillation strength will show near-zero values across the grid space, which compared to Figure 5b, has been dampened drastically.

By then merging the predictive neighborhood action sub-model and the dynamic action preference sub-model, we observe how the neighborhood prediction portion seems to

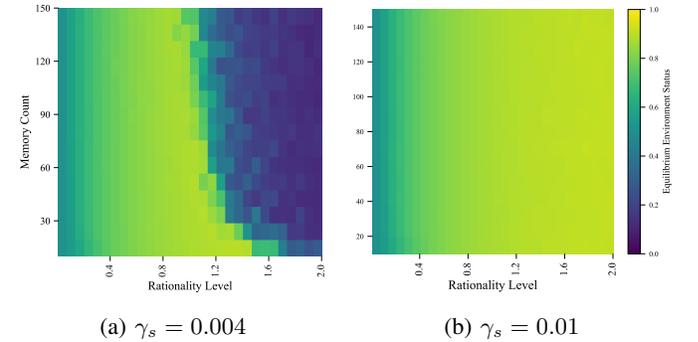
reinforce oscillation behavior with increasing rationality in Figure 8a when compared to Figure 7b, which causes the model with constant social pressure weights to stop moving towards a 0 equilibrium in Figure 8b when compared to Figure 7a.



(a) Random social pressure weight for all agents (b) 1.0 social pressure weight for all agents

Fig. 8: Environmental status at time step 1000 for an dynamic action preference, neighborhood action prediction, and single radius neighborhood

Figure 9 provides some evidence of criticality at $\gamma_s \approx 0.004$. The phase plot in Figure 9a shows three distinct regimes when varying memory count mem and rationality λ : when λ is very low (< 0.5), the equilibrium environmental state stays around neutral ($n^* \approx 0.5$) regardless of memory; when λ is high but memory is short ($mem < 40$), the environment collapses ($n^* < 0.2$); and when both λ and m exceed their respective thresholds (λ large and $mem > 60$), the system rapidly transitions to a healthy state ($n^* > 0.8$), with the critical boundary between collapse and recovery tracing a downward-sloping curve in the (mem, λ) plane.



(a) $\gamma_s = 0.004$

(b) $\gamma_s = 0.01$

Fig. 9: Phase diagrams of equilibrium environment state when varying memory count (mem) of agents and rationality λ . Figures (a) and (b) compares equilibrium environment state when agents adapt their action preference at slow rate ($\gamma_s = 0.004$) to when agents adapt at a faster rate ($\gamma_s = 0.01$).

Notably, small changes in either parameter near this boundary produce large shifts in the long-run environmental outcome, reflecting a tipping-point behavior. Additionally, the steepness of the transition becomes more pronounced at higher memory lengths, underlining how accumulated social information sharpens the collective move toward cooperation.

At higher action preference adaptation rate ($\gamma_s = 0.01$), we observe (Figure 9b) that the agents prefer to cooperate and choose the action that prevents the environment from collapsing, even for varying values of memory count and rationality. Phase diagrams for intermediate values of γ_s show this transition but the figures have not been included here for the sake of brevity.

B. Behavior at Critical Phase Transition

Phase plots offer a qualitative understanding of the model's behavior by exploring regions of the parameter space that exhibit large-scale dynamics. These visualizations help reveal unique system behaviors, such as the phase transitions and emergent clustering patterns observed in the dynamic action preference experiments.

For example, in Figure 7b, a distinct phase transition consistently appears which coincides with observable clustering behavior in Figure 7d. To locate the point at which this critical transition occurs, we can identify the steepest gradient from a heatmap slice in Figure 10a, followed by an analysis of the distribution of cluster counts near this critical point, as shown in Figure 10b.

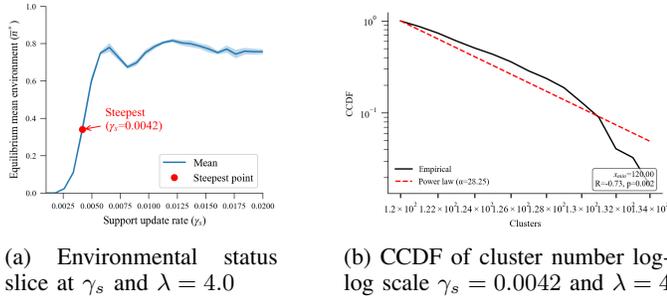


Fig. 10: Dynamic action preference, no neighborhood action prediction, single radius neighborhood, and random social pressure weights for all agents

In this model configuration featuring dynamic action preference without neighborhood action prediction, the cluster count distribution appears visually distinct from the fitted power-law with exponent $\alpha = 28.25$ (see section C). However, statistical analysis supports that the power-law is a good fit for values above $x_{\min} = 120$, where the distribution passes a significance test with $p = 0.002$. This visual mismatch is expected in empirical power-law distributions, as they often deviate at lower values due to finite-size effects, truncation, or under-sampling. The portion of the distribution that does conform to power-law scaling suggests that the system self-organizes near the critical regime. This supports the presence of scale-free clustering behavior, as seen in Figure 7b and Figure 10a.

The phase plot in Figure 8a adds linear neighborhood action prediction with a resulting distinct phase transition similar to Figure 7b. By once again taking the slice in Figure 11a and identifying the critical point $\lambda = 0.0048$, we get unique clustering behavior that wasn't observable in the previous experiments — larger zero-environment red clusters continue to exist alongside oscillations in the surrounding non-red cells,

but now small stable clusters of green high-environment can be observed in the grid as well.

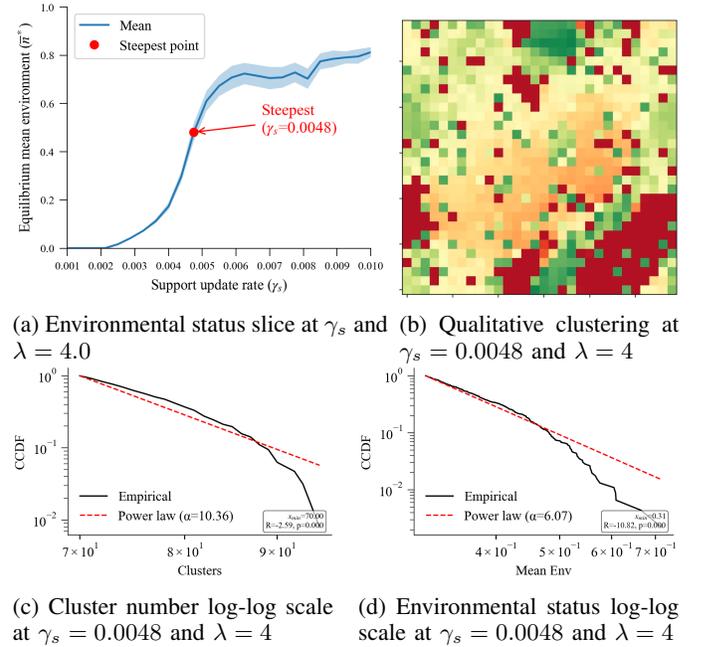


Fig. 11: Dynamic action preference, linear neighborhood action prediction, single radius neighborhood, and random social pressure weights for all agents

If we once again look at the cluster number around the critical point $\lambda = 0.0048$ and $\gamma = 4.0$ in Figure 11c, we observe possible scale-free behavior from $x_{\min} = 70$ onward with a statistically significant p -score approaching 0.00. By plotting another distribution over environmental status, we see an even closer fitting p -score approaching 0.00 with $x_{\min} = 0.31$.

C. Comparison with the mean-field model

Our mean-field analysis indicated that — if the underlying assumptions were well-founded — convergence to stable equilibria is guaranteed and well-defined solely by the expected peer pressure coefficient (w) and rationality parameter (λ). The results presented earlier in this section show a clear deviation from this indication.

First, while the PAWN sensitivity analysis (Figure 3) revealed a moderate primary effect of rationality on mean equilibrium environment state, it also indicated similar-level influences from the system size, as well as agent and environment adaptation rates. The latter two are of particular interest. Since equilibria are defined by zero derivatives in environment and agent actions, an equilibrium state which is sensitive to these parameters is indicative of both path-dependence in the steady-state, and the potential for more equilibria than possible in the mean-field model.

This discrepancy is (at least partially) explained by the clustering results observed in Figures 10 and 11, solidifying regions of consistent behaviour, locked in place by *locally-experienced* social norms. Such behaviour is not possible in

the mean-field model given its base assumptions. Overall, the results in this section allude to the presence of complex local behaviour arising from the interactions between social norms and the coupled human-dynamic system.

VI. CONCLUSION

This model attempts to bridge a research gap [2], [3] observed in the study of feedback between environmental status and agent actions, with the aim of better representing how diverse individual decision rules, social influences, and dynamic risk perceptions jointly determine collective environmental outcomes.

The model operationalizes the Theory of Planned Behaviour, balancing environmental preferences, social pressures, and perceived action feasibility through the use of dynamic risk-perception thresholds, predictive neighborhood influences, and heterogeneity in agent behaviors.

Through comprehensive sensitivity analysis using the PAWN method, it was found that memory size primarily influences mean environment and peak frequency, while mean action and pluralistic ignorance are most affected by recovery rate. Cluster count and dominant frequency power are significantly shaped by the adaptation rate, whereas spatial parameters such as neighborhood radius and grid length show minimal influence in any measured outcome.

Phase plot analysis revealed complex and dynamic system behaviors, including noisy periodic behavior in the base model which was dampened by the addition of the linear neighborhood predictor. The dynamic action preference revealed parameter-dependent oscillations and transitions between various system states, from stable to unstable oscillations and convergence to both zero and non-zero equilibria within separate distinct zones.

The research also provides evidence of criticality and tipping-point behavior where the model sees neutral equilibrium, environmental collapse, and a rapid transition to a healthy state when both rationality and memory exceeded certain thresholds. Around these critical points small changes lead to large shift in system-wide behavior, while simulations around these points exhibit possible scale-free clustering behavior.

Every component of this model is an attempt to define behaviors that are observed in the real world. Where individuals become complacent with a good environment and pessimistic in a bad one. Where friends and family influence the desire to act to the detriment or benefit of ourselves. An Agent-Based Model then provides a broad view of the chaos created by these interacting systems and individual preferences. But even in this chaos do unique structures form, creating equilibrium, oscillation, stochasticity, and clustering from the simple act of individual interaction.

REFERENCES

- [1] Council of the European Union. Fit for 55, 2025.
- [2] Oscar Kraan, Steven Dalderop, Gert Jan Kramer, and Igor Nikolic. Jumping to a better world: An agent-based exploration of criticality in low-carbon energy transitions. 47:156–165.
- [3] Andrew R. Tilman, Joshua B. Plotkin, and Erol Akçay. Evolutionary games with environmental feedbacks. 11(1):915. Publisher: Nature Publishing Group.
- [4] Donald L DeAngelis and Stephanie G Diaz. Decision-making in agent-based modeling: A current review and future prospectus. *Frontiers in Ecology and Evolution*, 6:237, 2019.
- [5] Wander Jager and Andreas Ernst. Introduction of the special issue: "social simulation in environmental psychology". *Journal of Environmental Psychology*, 52:114–118, 2017.
- [6] Jule Schulze, Birgit Müller, Jürgen Groeneveld, and Volker Grimm. Agent-based modelling of social-ecological systems: achievements, challenges, and a way forward. *Journal of Artificial Societies and Social Simulation*, 20(2), 2017.
- [7] Icek Ajzen. The theory of planned behavior. 50(2):179–211.
- [8] Evelin Ribeiro-Rodrigues and Ana Paula Bortoleto. A systematic review of agent-based modeling and simulation applications for analyzing pro-environmental behaviors. 47:343–362.
- [9] Laila Refiana Said, Fifi Swandari, Sufi Jikrillah, Sausan Sausan, and Fathia Azizah. Advancing self-social engineering in tourism-related environmental management: Integrating environmental psychology, planned behavior, and norm activation theories. 6(1):6. Number: 1 Publisher: Multidisciplinary Digital Publishing Institute.
- [10] George L. W. Perry, Sarah J. Richardson, Niki Harré, Dave Hodges, Phil O'B Lyver, Fleur J. F. Maseyk, Riki Taylor, Jacqui H. Todd, Jason M. Tylianakis, Johanna Yletyinen, and Ann Brower. Evaluating the role of social norms in fostering pro-environmental behaviors. 9. Publisher: Frontiers.
- [11] William A. Brock and Steven N. Durlauf. A formal model of theory choice in science. 14(1):113–130.
- [12] Sara M. Constantino, Gregg Sparkman, Gordon T. Kraft-Todd, Cristina Bicchieri, Damon Centola, Bettina Shell-Duncan, Sonja Vogt, and Elke U. Weber. Scaling up change: A critical review and practical guide to harnessing social norms for climate action. 23(2):50–97. Publisher: SAGE Publications Inc.
- [13] Zack Dörner. A behavioral rebound effect. 98:102257.
- [14] Alessandro Tavoni, Maja Schlüter, and Simon Levin. The survival of the conformist: Social pressure and renewable resource management. 299:152–161.
- [15] Joanna Aleksandra Kreft. The agent-based model of pluralistic ignorance, sincerity, conformity and truth. Accepted: 2024-09-04T04:05:16Z Publisher: UiT Norges arktiske universitet.
- [16] Alessandra F. Lütz and Lucas Wardil. The evolution of pluralistic ignorance. 647:129920.
- [17] Laura Pasca and Lucía Poggio. Biased perception of the environmental impact of everyday behaviors. 163(4):515–521. Publisher: Taylor & Francis Ltd.
- [18] Anne M. van Valkengoed, Goda Perlaviciute, and Linda Steg. From believing in climate change to adapting to climate change: The role of risk perception and efficacy beliefs. 44(3):553–565. Publisher: Wiley-Blackwell.
- [19] Van R. Haden, Meredith T. Niles, Mark Lubell, Joshua Perlman, and Louise E. Jackson. Global and local concerns: What attitudes and beliefs motivate farmers to mitigate and adapt to climate change? 7(12):e52882. Publisher: Public Library of Science.
- [20] Rocco Scolozzi, Anna Scolobig, and Marco Borga. Public support for flood risk management: Insights from an Italian alpine survey using systems thinking. 5(1):3. Number: 1 Publisher: Multidisciplinary Digital Publishing Institute.
- [21] Nurit Carmi, , and Shaul Kimhi. Further than the eye can see: Psychological distance and perception of environmental threats. 21(8):2239–2257. Publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/10807039.2015.1046419>.
- [22] Mete Sefa Uysal, Nuria Martinez, and Sara Vestergren. The horror of today and the terror of tomorrow: The role of future existential risks and present-day political risks in climate activism. 64(1):e12821. <https://bpspsychub.onlinelibrary.wiley.com/doi/pdf/10.1111/bjso.12821>.
- [23] Susie Wang, Mark J. Hurlstone, Zoe Leviston, Iain Walker, and Carmen Lawrence. Climate change from a distance: An analysis of construal level and psychological distance from climate change. 10. Publisher: Frontiers.

- [24] Roberta Maiella, Pasquale La Malva, Daniela Marchetti, Elena Pomarico, Adolfo Di Crosta, Rocco Palumbo, Luca Cetara, Alberto Di Domenico, and Maria Cristina Verrocchio. The psychological distance and climate change: A systematic review on the mitigation and adaptation behaviors. 11. Publisher: Frontiers.
- [25] Anne M. van Valkengoed, Linda Steg, and Goda Perlaviciute. The psychological distance of climate change is overestimated. 6(4):362–391.
- [26] Renée J. Johnson and Michael J. Scicchitano. Uncertainty, risk, trust, and information: Public perceptions of environmental issues and willingness to take action. 28(3):633–647.
- [27] Osman M. Jama, Abdishakur W. Diriye, and Abdulhakim M. Abdi. Understanding young people's perception toward forestation as a strategy to mitigate climate change in a post-conflict developing country. 25(6):4787–4811.
- [28] Tobias Kube, Jasmin Huhn, and Claudia Menzel. Optimistic bias in updating beliefs about climate change longitudinally predicts low pro-environmental behaviour. 64(3):e12905. <https://bpspsychub.onlinelibrary.wiley.com/doi/pdf/10.1111/bjso.12905>.
- [29] Jeroen C. J. M. van den Bergh, Ivan Savin, and Stefan Drews. Evolution of opinions in the growth-vs-environment debate: Extended replicator dynamics. 109:84–100.
- [30] Isaiah Farahbakhsh, Chris T. Bauch, and Madhur Anand. Tipping points in coupled human–environment system models: a review. 15(4):947–967.
- [31] Tarik Abou-Chadi and Mark A. Kayser. It's not easy being green: Why voters punish parties for environmental policies during economic downturns. 45:201–207.
- [32] Frances C. Moore, Katherine Lacasse, Katharine J. Mach, Yoon Ah Shin, Louis J. Gross, and Brian Beckage. Determinants of emissions pathways in the coupled climate–social system. 603(7899):103–111. Publisher: Nature Publishing Group.
- [33] Ana S. L. Rodrigues, Sophie Monsarrat, Anne Charpentier, Thomas M. Brooks, Michael Hoffmann, Randall Reeves, Maria L. D. Palomares, and Samuel T. Turvey. Unshifting the baseline: a framework for documenting historical population changes and assessing long-term anthropogenic impacts. 374(1788):1–9. Publisher: Royal Society.
- [34] Masashi Soga and Kevin J. Gaston. Global synthesis indicates widespread occurrence of shifting baseline syndrome. 74(10):686–694.
- [35] Frances C. Moore, Nick Obradovich, Flavio Lehner, and Patrick Baylis. Rapidly declining remarkability of temperature anomalies may obscure public perception of climate change. 116(11):4905–4910. Publisher: National Academy of Sciences.
- [36] Loren McClenachan, Ryuunosuke Matsuura, Payal Shah, and Sahan T. M. Dissanayake. Shifted baselines reduce willingness to pay for conservation. 5. Publisher: Frontiers.
- [37] Renee Zahnow, Ali Rad Yousefnia, Mahnoosh Hassankhani, and Ali Cheshmehzangi. Climate change inequalities: A systematic review of disparities in access to mitigation and adaptation measures. 165:104021.
- [38] Jonathan R. Olsen, Claire L. Niedzwiedz, Ruth Dundas, Jala Rizeq, Gergő Baranyi, Srinivasa Vittal Katikireddi, and Jill Pell. Linkages between socioeconomic inequalities, pro-environmental behaviours, climate change concerns and experiences, and wellbeing outcomes in England. 11(1):2500182.
- [39] Tao Wang, Feifei Liu, Shasha Xiong, and Kui Wang. Getting licence or keeping consistent? the role of moral identity in subsequent pro-environmental behaviours. 28(1):e12670. <https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajsp.12670>.
- [40] Heather Barnes Truelove, Amanda R. Carrico, Elke U. Weber, Kaitlin Toner Raimi, and Michael P. Vandenbergh. Positive and negative spillover of pro-environmental behavior: An integrative review and theoretical framework. 29:127–138.
- [41] Maedeh Gholamzadehmir, Paul Sparks, and Tom Farsides. Moral licensing, moral cleansing and pro-environmental behaviour: The moderating role of pro-environmental attitudes. 65:101334.
- [42] Jan Urban, Štěpán Bahník, and Markéta Braun Kohlová. Green consumption does not make people cheat: Three attempts to replicate moral licensing effect due to pro-environmental behavior. 63:139–147.
- [43] Andreas Nilsson, Bergquist, Magnus, and Wesley P. Schultz. Spillover effects in environmental behaviors, across time and context: a review and research agenda. 23(4):573–589. Publisher: Routledge. eprint: <https://doi.org/10.1080/13504622.2016.1250148>.
- [44] Adrian Brügger and Bettina Höchli. The role of attitude strength in behavioral spillover: Attitude matters—but not necessarily as a moderator. 10. Publisher: Frontiers.
- [45] B. L. Turner, Pamela A. Matson, James J. McCarthy, Robert W. Corell, Lindsey Christensen, Noelle Eckley, Grete K. Hovelsrud-Broda, Jeanne X. Kasperson, Roger E. Kasperson, Amy Luers, Marybeth L. Martello, Svein Mathiesen, Rosamond Naylor, Colin Polsky, Alexander Pulsipher, Andrew Schiller, Henrik Selin, and Nicholas Tyler. Illustrating the coupled human–environment system for vulnerability analysis: Three case studies. 100(14):8080–8085. Publisher: Proceedings of the National Academy of Sciences.
- [46] Jianguo Liu, Thomas Dietz, Stephen R. Carpenter, Carl Folke, Marina Alberti, Charles L. Redman, Stephen H. Schneider, Elinor Ostrom, Alice N. Pell, Jane Lubchenco, William W. Taylor, Zhiyun Ouyang, Peter Deadman, Timothy Kratz, and William Provencher. Coupled human and natural systems. 36(8):639–649. Publisher: Royal Swedish Academy of Sciences.
- [47] Chris T. Bauch, Ram Sigdel, Joe Pharaon, and Madhur Anand. Early warning signals of regime shifts in coupled human–environment systems. 113(51):14560–14567. Publisher: Proceedings of the National Academy of Sciences.
- [48] Brian Beckage, Louis J. Gross, Katherine Lacasse, Eric Carr, Sara S. Metcalf, Jonathan M. Winter, Peter D. Howe, Nina Fefferman, Travis Franck, Asim Zia, Ann Kinzig, and Forrest M. Hoffman. Linking models of human behaviour and climate alters projected climate change. 8(1):79–84. Publisher: Nature Publishing Group.
- [49] Isaiah Farahbakhsh, Chris T. Bauch, and Madhur Anand. Modelling coupled human–environment complexity for the future of the biosphere: strengths, gaps and promising directions. 377(1857):20210382. Publisher: Royal Society.
- [50] Joshua S. Weitz, Ceyhan Eksin, Keith Paarporn, Sam P. Brown, and William C. Ratcliff. An oscillating tragedy of the commons in replicator dynamics with game–environment feedback. 113(47).
- [51] Andrew R. Tilman, Vítor V. Vasconcelos, Erol Akçay, and Joshua B. Plotkin. The evolution of forecasting for decision-making in dynamic environments. 2(4):26339137231221726. Publisher: SAGE Publications.
- [52] Juana Castro, Stefan Drews, Filippos Exadaktylos, Joël Foramitti, Franziska Klein, Théo Konc, Ivan Savin, and Jeroen van den Bergh. A review of agent-based modeling of climate-energy policy. 11(4):e647.
- [53] Ivan Savin, Felix Creutzig, Tatiana Filatova, Joël Foramitti, Théo Konc, Leila Niamir, Karolina Safarzynska, and Jeroen van den Bergh. Agent-based modeling to integrate elements from different disciplines for ambitious climate policy. 14(2):e811.
- [54] Ram P. Sigdel, Madhur Anand, and Chris T. Bauch. Competition between injunctive social norms and conservation priorities gives rise to complex dynamics in a model of forest growth and opinion dynamics. 432:132–140.
- [55] Yoon Ah Shin, Sara M. Constantino, Brian Beckage, and Katherine Lacasse. Climate change and opinion dynamics models: Linking individual, social, and institutional level changes. 64:101528.
- [56] Fang Yan, Xiaojie Chen, Zhipeng Qiu, and Attila Szolnoki. Co-operator driven oscillation in a time-delayed feedback-evolving game. 23(5):053017. Publisher: IOP Publishing.
- [57] Fatima Abdallah, Shadi Basurra, and Mohamed Medhat Gaber. An agent-based collective model to simulate peer pressure effect on energy consumption. In Ngoc Thanh Nguyen, Elias Pimenidis, Zaheer Khan, and Bogdan Trawiński, editors, *Computational Collective Intelligence*, volume 11055, pages 283–296. Springer International Publishing. Series Title: Lecture Notes in Computer Science.
- [58] Ezgi Topuz and Gönenç Yücel. Analyzing the emergence and dynamics of pluralistic ignorance with agent-based models. In Corinna Elsenbroich and Harko Verhagen, editors, *Advances in Social Simulation*, pages 423–434. Springer Nature Switzerland.
- [59] Dhruv Mittal, Sara M Constantino, and Vítor V Vasconcelos. Anti-conformists catalyze societal transitions and facilitate the expression of evolving preferences. 3(8):pgae302.
- [60] Marten Scheffer, Frances Westley, and William Brock. Slow response of societies to new problems: Causes and costs. 6(5):493–502.
- [61] Marten Scheffer. *Critical Transitions in Nature and Society*. Princeton University Press.
- [62] Andrea Saltelli. Sensitivity analysis for importance assessment. *Risk analysis*, 22(3):579–590, 2002.
- [63] Francesca Pianosi and Thorsten Wagener. A simple and efficient method for global sensitivity analysis based on cumulative distribution functions. 67:1–11.
- [64] Jeff Alstott, Ed Bullmore, and Dietmar Plenz. powerlaw: A python package for analysis of heavy-tailed distributions. *PLoS ONE*, 9(1):e85777, 2014.

APPENDIX A

ODD+D IMPLEMENTATION DETAILS

A. Overview

1) *Purpose and Entities*: The model explores how individual decisions to adopt pro-environmental behaviors are shaped by environmental feedback and social norms, where agents are arranged on a 2-D grid and repeatedly choose between positive and negative environmental actions. Each agent has an adaptive behavioral preference, observes a local environmental state, and is influenced by their prediction of neighbors' past actions.

2) *Process and scheduling*: Simulations run in discrete time-steps that each follow these specific steps:

- (a) *Environment Update*: Agents' environments evolve based on prior actions, using a user-definable function.
- (b) *Behavioral Adaptation*: Agents update their weight placed on positive or negative action in response to environmental conditions. Willingness to take positive action changes more strongly when the environment is neither fully degraded nor fully healthy.
- (c) *Prediction of Neighbors' Actions*: When enabled, agents estimate future neighbor behavior using linear or logistic prediction and a specified amount of memory of neighbor actions.
- (d) *Decision*: Agents select actions probabilistically based on their updated willingness to take positive or negative action and predicted social norms, as defined in the logit equation.

B. Design concepts

1) *Theoretical Basis and Decision Making*: The model builds on theories of bounded rationality and norm-driven behavior, making decisions using a probabilistic logit model based on each agent's support and the predicted behavior of their neighbors.

2) *Agent Learning, Sensing, and Predicting*: Agents sense their environment and recent neighbor actions, adjusting their tendency for action depending on environmental quality, taking more positive action in moderately degraded settings and less at environmental extremes. Agents also sense their neighbors' actions and retain a specified memory of those actions to predict overall neighborhood behavior using recent trends.

3) *Interaction*: Agents interact with each other in a Moore neighborhood. The neighbors' actions are taken into consideration when an agent is trying to maximize its expected utility (as shown in Equation 1). An agent also directly interacts with the environment in it's assigned location on the lattice by choosing an action that is dependent on the current state of the environment.

4) *Heterogeneity*: Each agent in the model has an independent preference for climate action. An agent also adapts its behaviour dependent on the state of the environment at its position in the lattice. In experiments with heterogeneous social influence, each agent's peer-pressure coefficient is initialized by sampling independently from a Uniform(0, 1) distribution, thereby seeding behavioral diversity.

5) *Stochasticity*: Stochasticity in the model is introduced primarily through agents' action choices, which are drawn at each time step from a logit (softmax) distribution (Equation 2) over their computed utilities. All results are obtained via multiple Monte Carlo replicates, each with a fresh random seed, to ensure that observed patterns reflect the variability induced by these stochastic processes.

6) *Collectives*: All agents are part of a neighborhood, and the action level of the neighborhood affects the environmental state of a given cell. Clusters of healthy or degraded environment form as a result of agent interaction with its neighbors and its environment. Cluster formation is an emergent property of the model.

7) *Observation*: Time-series data for the environmental state and agent action is collected for many different parameter sets to analyze model behaviour. Some of the recorded metrics are: mean environment and action, pluralistic ignorance, spatial cluster counts and size distributions, and oscillation characteristics via Fourier analysis. These measurements are used to validate the observed model dynamics with theoretical expectations.

C. Details

1) *Parameters*: Experimental parameters listed in Table II.

Parameter	Description	Values
N	Number of agents	30×30
λ	Agent rationality	[0, 6]
γ_s	Action preference rate of change	[0.001, 0.02]
b_1	Weight on environmental concern	Uniform 1.0
b_2	Weight on social norms	Random [0.0, 1.0] and Uniform 1.0
$s(t)$	Initial support for cooperation	Uniform 1.0
T	Simulation duration (timesteps)	1000
m	Memory length (steps)	10
r	Radius of local neighborhood	1 or Mean-field
<code>env_up_fn</code>	Environmental update function	linear or exponential
<code>env_adap</code>	Calculation of environment change	linear or adaptive
<code>prediction</code>	Neighbor behavior prediction	None or linear
<code>repeats</code>	Simulation repeats per condition	10 – 100
<code>moore</code>	Moore Neighborhood Option	True

TABLE II: Model Parameters Used in Simulation Experiments

2) *Implementation and Initialization*: The model is implemented as both an object-oriented and vectorized agent-based simulation in Python, using a 2-D spatial grid where each cell represents a single agent who continuously chooses between improving or impairing their environment. The object-oriented model emphasizes modularity and ease-of-use, while the vector-based model emphasizes efficiency for large-scale experimentation — the vectorized model contains all of the most recent changes, so focus will be placed on that.

As opposed to the object-oriented method, the vectorized model encodes agents as rows in a structured array instead of as individual objects, where each array stores one attribute of the individual agent across multiple time steps, including actions, environmental states, or parameter values. This structure allows for vectorized operations, in which entire populations are updated in a single step using high-performance libraries like NumPy and SciPy.

3) Submodels:

- Environmental Update Dynamics*: Either linear or piecewise exponential.
- Environmental Action Dynamics*: Driven by nonlinear response to environmental quality.
- Local Action Prediction*: Predict neighbor actions using linear or logistic regression.
- Decision Model*: Logit-based decision making between positive and negative action.

4) *Experiments*: Systematic experiments were conducted on every combination of the parameters specified in Table II and plotting γ_s and λ on heatmaps of different output values, including mean environmental status, number and size of clusters, and Fourier analysis of cyclic behaviors.

General behavior of the system was identified with the heatmaps and then isolated into slices of specific continuous parameter values to find the steepest critical point for qualitative phase transitions. Using the specific parameter value identified at the critical point, plots of model outcomes over time were created to observe behavior before, after, and at the critical point. These results can then be compared to the mean-field approximation to observe differences caused by the model implementation.

APPENDIX B MEAN-FIELD DERIVATIONS

A. Expected agent action

Consider the action probability as defined in Equation 2. By expanding the representative utility and simplifying, we can eliminate any terms not dependent on a (the specific action). For clarity, we omit the t superscripts on a_i^t and s_i^t :

$$V_i(a) = (1 - w_i)[2 + (s_i - 2)a] - w(a - \bar{a}_i)^2 \quad (8)$$

$$= (1 - w_i)(2 + s_i a - 2a) - w(1 - 2a\bar{a}_i + \bar{a}_i^2) \quad (9)$$

$$= [2(1 - w_i) - w(1 + \bar{a}_i^2)] \quad (10)$$

$$+ a[(1 - w_i)(s_i - 2) + 2w\bar{a}_i] \quad (11)$$

Taking $z_i = (1 - w_i)(s_i - 2) + 2w\bar{a}_i$, we derive agent i 's expected action:

$$\mathbb{E}[a_i] = \mathbb{P}(a_i = C) - \mathbb{P}(a_i = D) \quad (12)$$

$$= \frac{\exp(\lambda z_i) - \exp(-\lambda z_i)}{\exp(\lambda z_i) + \exp(-\lambda z_i)} \quad (13)$$

$$= \tanh(\lambda(1 - w_i)(s_i - 2) + 2\lambda w\bar{a}_i) \quad (14)$$

B. Expected mean action

To obtain an expression for the population-level expected action, we first make the following simplifying assumptions:

- The heterogeneous w_i attributes assume their expected value, $w_i = w$.
- Agents' adapt sufficiently fast, such that their action probability — and thus action preferences — are stationary.
- For any pair of agents, their action probabilities are independent (the mean-field assumption).

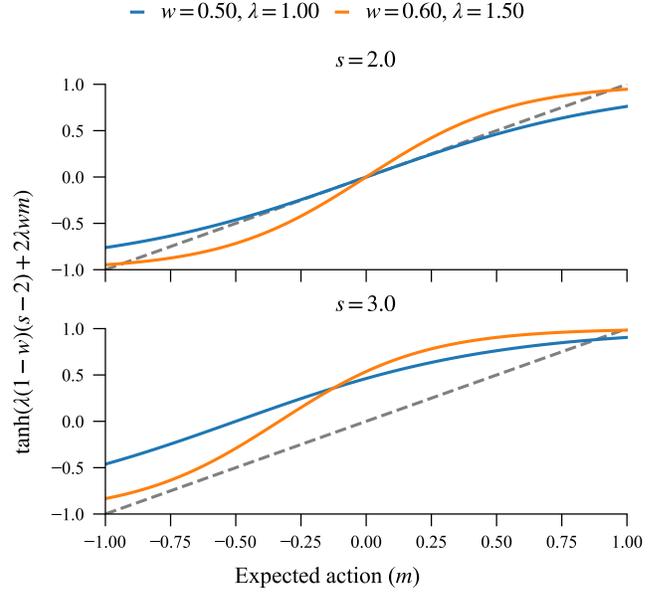


Fig. 12: Graphical solutions (fixed points) of the average action mean-field equation for two levels of initial support for cooperation (top) $s = 2.0$ and (bottom) $s = 3.0$ for varying levels of social-norm coefficient (w) and rationality (λ).

Taken together, these assumptions imply that $\bar{a}_i = m$ for each agent $i \in N$, where m is the population-level expected action, and $w_i = w$, $s_i = s$. Thus, the expected mean action is:

$$m = \mathbb{E} \left[\frac{1}{k} \sum_{i=1}^k a_i \right] \quad (15)$$

$$= \frac{1}{k} (\mathbb{E}[a_1] + \dots + \mathbb{E}[a_k]) \quad (16)$$

$$= \frac{1}{k} \cdot k \cdot \mathbb{E}[a_i] \quad (17)$$

$$= \tanh(\lambda(1 - w)(s - 2) + 2\lambda w m) \quad (18)$$

The fixed-point equation is illustrated with varying w, λ, s in Figure 12. Equation 18 has between 1 and 3 solutions. We now characterize the necessary and sufficient conditions for three solutions to exist. Rewriting Equation 18 as a fixed-point problem $f(m)$,

$$f(m) = g(m) - m \quad (19)$$

$$g(m) = \tanh(\lambda(1 - w)(s - 2) + 2\lambda w m) \quad (20)$$

we observe that $f(m) = 0$ can have three solutions, only if $g'(m) > 1$ for some $m \in [-1, 1]$. This condition is satisfied when:

$$1 < g'(m) = \text{sech}^2(\lambda z) \cdot 2\lambda w \leq 2\lambda w \quad (21)$$

Thus for three solutions to exist, it is necessary that $\frac{1}{2\lambda w} < 1$, i.e., individual rationality and the influence of social norms must be sufficiently high. Now observe that for multiple solutions to exist, it is sufficient to show that:

- f has two turning points within the domain $m \in [-1, 1]$, and that

2) These correspond to expected actions m_1^*, m_2^* which lie below and above the diagonal, respectively.

With $f'(m) = 2\lambda w \cdot \text{sech}^2(\lambda z) - 1$, the turning points are given by:

$$m^* = -\frac{(1-w)(s-2)}{2w} \pm \frac{1}{2\lambda w} \text{arccosh}(\sqrt{2\lambda w}) \quad (22)$$

C. Mean-field dynamics

We derive the temporal dynamics of m by taking the time derivative of Equation 18:

$$\frac{dm}{dt} = \frac{(1-w)\lambda \text{sech}^2(\lambda z)}{1-2\lambda w \text{sech}^2(\lambda z)} \cdot \frac{ds}{dt} \quad (23)$$

Where the mean-field $\frac{ds}{dt}$ is defined by:

$$\frac{1}{\gamma_s} \frac{ds}{dt} = \beta_s^+ \cdot \sigma(n)(4-s) - \beta_s^- \cdot (1-\sigma(n))s \quad (24)$$

Observe that as the denominator in Equation 23 approaches 0, the derivative diverges. This occurs precisely at the identified turning points of the fixed-point function f . Thus when $\frac{1}{2\lambda w} < 1$, the mean-field model predicts increasing attraction to the two stable fixed points as they are approached.

Equilibria in the mean-field dynamical system are characterised by mean environment n^* and action m^* such that $\frac{dn}{dt} = \frac{dm}{dt} = 0$. The temporal environment derivative in the mean-field model is given by:

$$\frac{1}{\gamma_n} \frac{dn}{dt} = \beta_n^+ \cdot (1-n)\mathbb{P}(C) - \beta_n^- \cdot n\mathbb{P}(D) \quad (25)$$

Thus the environment is stationary when:

$$0 = \beta_n^+ \cdot (1-n) - \beta_n^- \cdot n^t(1-\alpha_i^t) \quad (26)$$

$$\implies n^* = \frac{\beta_n^+ \mathbb{P}(C)}{\beta_n^+ \mathbb{P}(C) + \beta_n^- \mathbb{P}(D)} \quad (27)$$

From Equation 23, the expected action is stationary when:

$$(1-w)\lambda \text{sech}^2(\lambda z) \cdot \frac{ds}{dt} = 0 \quad (28)$$

Since $\text{sech} > 0$, we find that $\frac{dm}{dt} = 0$ when agents are unaffected by peer pressure, are completely irrational, or when the mean action preference is unchanging. The latter condition holds when:

$$s^* = \frac{4 \cdot \beta_s^+ \sigma(n)}{\beta_s^+ \sigma(n) + \beta_s^- (1-\sigma(n))} \quad (29)$$

Now suppose that n^* is such that $\frac{dn}{dt} = 0$. Solving for the corresponding $P_C^* = \mathbb{P}(C)$, we find:

$$P_C^* = \frac{\beta_n^- n^*}{\beta_n^- n^* + \beta_n^+ (1-n^*)} \quad (30)$$

However, from the logit model definition, we also have that:

$$P_C^* = \frac{\exp(\lambda z)}{\exp(\lambda z) + \exp(-\lambda z)} \quad (31)$$

$$= \frac{1}{1 + \exp(-2\lambda z)} \quad (32)$$

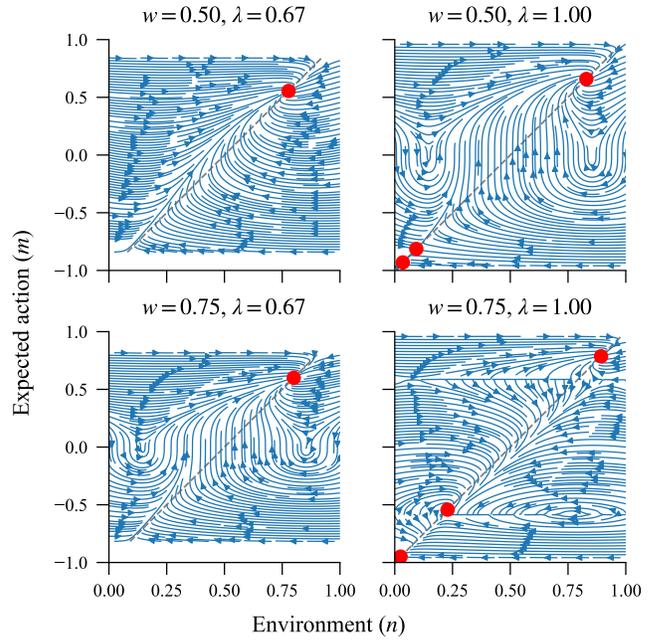


Fig. 13: Phase portraits of the mean field dynamics of our model for varying social-norm coefficient (w) and rationality level of agents (λ).

Thus we can rearrange for the corresponding steady-state action preference, s^* :

$$P_C^* = \frac{1}{1 + \exp(-2\lambda z)} \quad (33)$$

$$\frac{1}{P_C^*} = 1 + \exp(-2\lambda z) \quad (34)$$

$$-\frac{1}{2\lambda} \ln \left(\frac{1-P_C^*}{P_C^*} \right) = (1-w)(s^*-2) + 2w \cdot \overbrace{m}^{2P_C^*-1} \quad (35)$$

$$s^* = \frac{1}{1-w} \left[2(1-2wP_C^*) - \frac{1}{2\lambda} \ln \left(\frac{1-P_C^*}{P_C^*} \right) \right] \quad (36)$$

Since P_C^* depends on n^* , by equating Equations 29 and 36, we obtain a fixed-point problem, whose numerical solution defines the system equilibria (n^*, s^*) . Finally, we recover the equilibrium expected action as

$$m^* = P_C^* - P_D^* = 2P_C^* - 1 \quad (37)$$

Example dynamics are illustrated for varying parameters in Figure 13, with red circles identifying the equilibria, and the $\frac{dn}{dt} = 0$ nullcline depicted with a grey dashed line.

APPENDIX C POWER LAW FITTING

To determine whether certain quantities in the simulation follow a power-law distribution, we fit empirical distributions using the `powerlaw` Python package [64].

Power Law Model

A continuous power-law distribution is defined by the probability density function:

$$p(x) = Cx^{-\alpha}, \quad x \geq x_{\min}, \quad (38)$$

where $\alpha > 1$ is the scaling exponent, x_{\min} is the lower bound above which the power-law behavior holds, and C is a normalization constant.

Given a set of n observations $\{x_i\}$ such that $x_i \geq x_{\min}$, the scaling exponent α is estimated via the maximum likelihood estimator (MLE):

$$\alpha = 1 + n \left[\sum_{i=1}^n \ln \left(\frac{x_i}{x_{\min}} \right) \right]^{-1}. \quad (39)$$

Determining x_{\min}

To find the optimal value of x_{\min} , the Kolmogorov–Smirnov (KS) distance between the cumulative distribution function (CDF) and the fitted power-law model is computed for a range of x_{\min} . The value that minimizes this distance is selected:

$$x_{\min} = \arg \min_x \sup_{x \geq x_{\min}} |S(x) - P(x)|, \quad (40)$$

where $S(x)$ is the empirical CDF and $P(x)$ is the CDF of the fitted model.

Goodness-of-Fit Test and p -Value

To assess the power-law fit, a goodness-of-fit test is done by simulating the fitted model and computing their KS distances from the model. The p -value is then calculated as the fraction of these simulated distances that exceed the empirical KS distance:

$$p = \frac{\# \text{ simulations with } D_{\text{sim}} > D_{\text{emp}}}{\# \text{ simulations}}. \quad (41)$$

A p -value greater than 0.1 is generally taken to indicate that the power-law hypothesis cannot be rejected.

Model Comparison

To further validate the fit, we compute the log-likelihood ratio R between the power-law model and an alternative distribution (e.g., exponential). The corresponding p -value is calculated to determine if the observed difference is statistically significant:

$$R = \log \left(\frac{L_{\text{power law}}}{L_{\text{alternative}}} \right). \quad (42)$$

A positive R indicates that the data are more likely under the power-law model.