



The Strategic Dynamics of the Anthropocene: Bridging Social Dilemma Games and Empirical CO₂ Emissions with Probabilistic Programming

Ali Banihashemi¹, Hugo Storm², and Wolfram Barfuss^{1,2,3}

¹Center for Development Research (ZEF), University of Bonn, Bonn, Germany

²Institute for Food and Resource Economics (ILR), University of Bonn, Bonn, Germany

³Transdisciplinary Research Area Sustainable Futures, University of Bonn, Bonn, Germany

Correspondence: Ali Banihashemi (ali.banihashemi@uni-bonn.de) and Wolfram Barfuss (wbarfuss@uni-bonn.de)

Abstract. In the Anthropocene, human activities dominate planetary change, yet the strategic motivations behind this unsustainable behavior remain poorly understood. Game theory has been used to analyze the strategic structure of the Anthropocene; however, it lacks empirical grounding. Here, we bridge this gap by inferring the strategic dynamics underlying global CO₂ emissions. We model emissions as the outcome of stylized games, then we use Bayesian probabilistic programming to compare ten model configurations against historical data.

Our results indicate that the climate problem is more likely to be a free-riding problem, not a coordination failure. The inferred parameters capture a deep inequality in emissions, resonating with the Global North–South narrative and differentiated responsibilities. The projections align with RCP 6.0 (approximately 2.0–3.7°C warming by 2100; best estimate ~2.8°C).

Our results suggest that intervention toward transforming the payoff structure must be ordered: the first priority is controlling the dominant emitter through reducing gains from defection through enforceable mechanisms and technological transformation; the second is addressing fear-driven defection by the minor emitter through transfers that support vulnerable actors. Methodologically, we demonstrate that Bayesian probabilistic programming can bridge the gap between stylized game theoretic models and empirical data, offering a principled framework for inferring strategic incentives from observed trajectories while quantifying parameter uncertainty.

1 Introduction

The concept of Anthropocene signals both the scale of human impact and a shift in Earth system dynamics (Crutzen, 2002), in which human decisions and biophysical processes are tightly coupled (Steffen et al., 2015b). A hallmark of this transformation is the Great Acceleration: the exponential-like growth of socio-economic and biophysical trends beginning in the mid-20th century (Steffen et al., 2015a).

The consequences of anthropogenic change pose imminent risks to the resilience of the Earth system. Human activities destabilize the global carbon cycle and drive rapid loss of biosphere integrity, pushing the planetary system beyond its safe operating space (Richardson et al., 2023; Rockström et al., 2009). If current policies continue, the average global temperature is



projected to increase by approximately 3.1°C by 2100 (United Nations Environment Programme, 2024). This level increases the likelihood of crossing multiple critical tipping points, which can amplify warming via positive feedback (Steffen et al., 2018; Lenton et al., 2008). Even milder warming levels (1.5–2.0°C) may trigger some tipping elements (Armstrong McKay et al., 2022), and because of hysteresis effects, mitigation efforts may not be able to reverse triggered changes (Wunderling et al., 2021; Solomon et al., 2009).

Awareness of climate change risks has been growing, however, human responses remain insufficient indicating that the gap between knowledge and action is persistent. Early efforts to close the gap focused on communicating the physical science, assuming that lack of awareness was the primary barrier (Moser, 2010). However, even with the presence of broad public awareness and strong scientific consensus, the business-as-usual path mostly persists and emissions continue to rise (Friedlingstein et al., 2022). This suggests that the core problems are behavioral and structural rather than a lack of knowledge (Kollmuss and Agyeman, 2002; Stoknes, 2014). Psychological and structural barriers may explain why people often fail to act. At the behavioral level, cognitive constraints such as bounded rationality and short-term discounting encourage unsustainable behavior and the overlooking of future risks (Gifford, 2011; Frederick et al., 2002; Simon, 1955). Unconscious habitual routines (Ouellette and Wood, 1998) and moral disengagement also might undermine sustainable behavior (Markowitz and Shariff, 2012). At the structural level, personal interests and carbon dependence may foster a social resistance to change (Otto et al., 2020). Moreover, institutional reliance on fossil fuels leads to carbon lock-in, and hinders the adoption of cleaner energy resources (Unruh, 2000). In addition, inequality sabotages collective action even more, as asymmetries in historical responsibility, economic capacity, and vulnerability create divergent interests among nations (Klinsky et al., 2017; Paavola and Adger, 2006; Okereke, 2008). This unequal distribution of responsibilities, costs and benefits suggests that actors may perceive fundamentally different incentive structures, which undermines the formation of collaborative solutions (Lange et al., 2007).

Perhaps, these psychological, institutional, and political barriers lead to a common thread. All of them can shape the incentive structures that actors face when it comes to emission decisions. This raises a deeper diagnostic question about the underlying strategic nature of the climate problem. Game theory provides a natural framework for this question, although the answer remains contested (Aklin and Mildenberger, 2020). The climate collective-action problem is often presented as a tragedy of the commons (Hardin, 1968), but it has been given different strategic interpretations: as a Prisoner's Dilemma, in which defection is individually rational decision regardless of what others do, therefore, cooperation requires external enforcement (Barrett, 2003; Nordhaus, 2015); as a coordination game, in which actors would prefer to cooperate but fail to align on a joint strategy because of mistrust or misalignment in their expectations (Skyrms, 2004; Hale, 2020); and as a Chicken game, in which each actor prefers that others bear the cost of cooperation (DeCanio and Fremstad, 2013). Each diagnosis implies different policy levers, and adopting the wrong one risks misdirecting intervention which can be harmful (Carrozzo Magli and Manfredi, 2022). Thus, to design effective policy, we need to determine which social dilemma best explains the observed emissions trajectory.

Answering this question requires treating social dynamics as an endogenous component in Earth-system modeling (Donges et al., 2017). Given the scale of human impact, social processes and human responses are essential components of the Earth system that couple to the biophysical state through feedback (Donges et al., 2020). Nevertheless, the complexity of the whole World-Earth system creates a trade-off between tractability and realism in modeling. The Earth system is already highly



complex on its own, and human behavior, which is also challenging to formalize, adds another layer of difficulty on top. The more detailed a model becomes, the harder it is to use for explaining the phenomena it represents (Dutton and Starbuck, 1971).

60 Levins (1966) formalized this tension, showing that no model can simultaneously maximize generality, realism, and precision.

This tension is visible in Earth system models. Realistic models fit the data closely, but are often opaque and reproduce patterns without revealing the mechanisms behind them. General models offer clearer causal insight through abstraction, but their distance from the data makes them hard to validate.

Existing modeling approaches to Earth system dynamics, usually represent human behavior as either passively reactive or
65 fully optimal, rather than boundedly rational and adaptively strategic (Simon, 1955). Integrated assessment models (IAMs) couple economic and climate dynamics with enough precision for quantitative policy evaluation, but typically reduce human agency to a single abstract social planner choosing optimally (Keppo et al., 2021). Classical game-theoretic models rely on equilibria that assume optimal play (Lessmann et al., 2009). Agent-based models (ABMs) go further in capturing heterogeneous actors and emergent dynamics, but their agents typically follow fixed, context-specific heuristics, and the computational cost
70 of running them limits systematic exploration of the parameter space (Schulze et al., 2017). Dynamical systems approaches are powerful for identifying feedback, leverage points, and unintended consequences, but they tend to represent human responses (when represented at all) as reactive adjustments to system states rather than as strategic choices (Radosavljevic et al., 2023). However, many real sustainability challenges involve actors who anticipate and respond to each other, so the effect of any intervention depends on how others react.

75 Behavioral game theory has linked theoretical predictions to empirical observations in controlled laboratory experiments (McKelvey and Palfrey, 1995; Camerer, 2003; Crawford et al., 2013; Kuleshov and Schrijvers, 2015), and some studies have extended this to environmental settings such as pollution in shared waterways. To our knowledge, however, this empirical linkage has not been made for biophysical climate data. Learning and evolutionary dynamic models have been applied to climate games to account for adaptively strategic human behavior (Santos et al., 2012; Hilbe et al., 2013; Vasconcelos et al.,
80 2014; Barfuss et al., 2020), but their payoff structures often lack direct empirical grounding. Without calibration to data, such models cannot identify the effective strategic incentives that actually govern emissions, and so the models themselves risk remaining metaphorical. Empirically quantifying strategic parameters would move the analysis beyond stylized stories toward testable claims about the drivers of collective action (Kollock, 1998; Donges and Barfuss, 2017).

In this paper, we link learning dynamics in idealized models of the World-Earth system to historical CO₂ emissions. We
85 use Bayesian probabilistic programming (Storm et al., 2024; Krapu and Borsuk, 2019; Gelman et al., 2013) to estimate the parameters of social dilemma games while quantifying parameter uncertainty. Because the method returns a posterior distribution for each parameter rather than a point estimate, it lets us see directly which parameters are well-identified by the data and which are not (Grazzini et al., 2017; Hartig et al., 2011). Probabilistic programming of this kind has been applied in economics, ecology, and demography (Lamperti et al., 2018; Guerini and Moneta, 2017; Barde, 2016; van der Vaart et al., 2015; Bijak,
90 2021).

We apply this method to the modeling framework of learning dynamics in games (Bloembergen et al., 2015; Fudenberg and Levine, 2016; Barfuss and Mann, 2022). Learning and evolutionary dynamics sits between the modeling traditions mentioned



above and provides a useful middle ground. We focus on social dilemma games such as the Prisoner's Dilemma, Chicken, Stag Hunt, and Harmony, which differ in their strategic nature and imply different policy levers (Barrett, 2007; Lipman, 1986; 95 Skyrms, 2004). In our framework, two agents represent aggregate emitting blocs, each choosing a level of defection that represents prioritizing carbon-intensive activity for economic gain. Total CO₂ emissions are the observable outcome, modeled as a multiple of aggregate defection.

We do not claim that emitting blocs deliberately play a social dilemma. We ask whether the pattern of outcomes they collectively produce is better explained *as if* blocs learned to play under one strategic incentive structure rather than another, 100 a methodological stance with a long tradition in economics (Friedman, 1953). The inferred parameters, including the agents' motives, are effective values. They compress shifting motivations, evolving institutions, and the changing carbon footprints of human activity into the strategic quantities that best explain the observed trajectory.

Given the contested strategic architecture of the climate collective-action problem, we ask (i) how well each social dilemma game structure explains the data, and, given the central role of inequality in the climate problem, (ii) how important agent 105 heterogeneity is for the explanation.

We find that Prisoner's Dilemma and Chicken games generally fit the data well, whereas Stag Hunt and Harmony games fail. The inequality has a significant role in expressing the historical emissions and heterogeneous agents outperform. The identified policy directions are not themselves new; however, our contribution is the empirical identification of which game structure the data actually support, where leverage is concentrated, and consequently, which policy levers carry quantitative weight, rather 110 than relying on theory. We hope our study to be useful for connecting tractable models with empirical observations, fostering both approaches.

2 Methods

We calibrate the parameters of two-player social dilemma games with learning agents based on historical CO₂ emissions using Bayesian probabilistic programming. We first describe the process-based model of the dynamic game, then the probabilistic 115 programming procedure, and finally the parameter sensitivity analysis.

2.1 Process-based models

The process-based models are composed of a decision environment and a learning dynamic. First, we describe the decision environments of the social dilemma games. Second, we explain the learning dynamics that we used.

2.1.1 Social dilemma game environments

120 Two agents interact, each can choose to either cooperate, C, or defect, D. We investigate two parameterizations of a social dilemma, the widely used *cost-benefit* (CB) parameterization, and a more flexible *fear-greed* (FG) parameterization.

Cost-Benefit (CB) reparameterization. The CB form is a parameterization of the Prisoner's Dilemma that relates the free-riding problem to the benefits and costs of cooperation. Let $b > 0$ be the benefit an agent provides to the *other* agent by



cooperating, and let $c \geq 0$ be the private cost of doing so. We normalize the game by dividing all payoffs by benefit b and
 125 defining the cost-to-benefit ratio $r \equiv c/b$, with $0 \leq r \leq 1$. Thus, r is the sole model parameter of this parametrization.

Fear-Greed (FG) reparameterization. The CB reparameterization is prevalent and intuitive; however, it only covers the strategic incentives of a Prisoner’s Dilemma. In the context of climate games, other forms of a social dilemma, such as the Chicken and Stag-Hunt coordination game, have also been applied (Barrett, 2003).

Therefore, we also study another parametrization based on the *greed*, g , to exploit others, and the *fear*, f , of being exploited
 130 by others (Coombs, 1973; Macy and Flache, 2002). We set the payoffs for unilateral cooperation to 1 and mutual defection to 0. Consequently, the payoff for mutual cooperation becomes $1 - g$, and the payoff for being exploited becomes $-f$ (Fig. 1a). Note that this way, the CB parameterization is a subtype of this reparameterization where $f = g = r$.

The four strategic games correspond to the different combinations of signs of the fear and greed values (Fig. 1b). In the Prisoner’s Dilemma (FG_{pd}), agents are greedy to exploit others ($g > 0$) and fear being exploited by others ($f > 0$). In the
 135 Chicken game (FG_{ch}), agents are also greedy to exploit others ($g > 0$), but do not fear being exploited by other ($f < 0$), as they are better off with unilateral cooperation than mutual defection. In the Stag-Hunt coordination game (FG_{sh}), agents lack the temptation to exploit ($g < 0$), but remain fearful of being exploited ($f > 0$). And last, out of symmetry, in the Harmony game (FG_{ha}), agents are neither greedy to exploit ($g < 0$), nor fear being exploited by others ($f < 0$).

Thus, in total, we test five decision environments: {CB, FG_{pd}, FG_{ch}, FG_{sh}, FG_{ha}}, addressing our first question,
 140 how well each social dilemma game structure is able to explain the data.

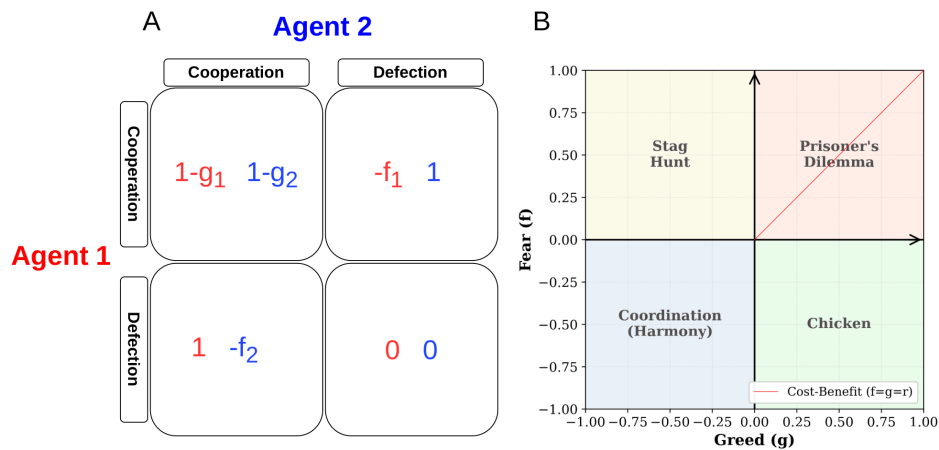


Figure 1. Fear-greed parameterization and the resulting strategic regimes. **(a)** The standard payoff matrix for a two-agent social dilemma. Payoffs are formulated such that mutual cooperation yields a Reward ($R = 1 - g_i$), mutual defection yields a Punishment ($P = 0$), unilateral defection yields Temptation ($T = 1$), and unilateral cooperation results in a Sucker’s payoff ($S = -f_i$). **(b)** The four strategic regimes defined by the parameter space of fear (f) and greed (g). The quadrants emerge from the relative ranking of payoffs: Prisoner’s Dilemma ($T > R > P > S$; $f > 0, g > 0$), Stag Hunt ($R > T > P > S$; $f > 0, g < 0$), Chicken ($T > R > S > P$; $f < 0, g > 0$), and Coordination/Harmony ($R > T > S > P$; $f < 0, g < 0$). The red line illustrates the CB (Cost-Benefit) parameterization where $r_i = f_i = g_i > 0$.



2.1.2 Learning dynamics

We use a general learning dynamic that emerges from perspectives in both evolutionary game theory (Tuyls and Parsons, 2007; García and Traulsen, 2025) and individual learning via reinforcements (Börgers and Sarin, 1997; Barfuss et al., 2019). Specifically, we use exponential replicator dynamics in discrete time (Hofbauer and Sigmund, 1998),

$$145 \quad X_{i,a}(t+1) = \frac{X_{i,a}(t) \exp\{\alpha_i R_{i,a}(t)\}}{\sum_{b \in \{C,D\}} X_{i,b}(t) \exp\{\alpha_i R_{i,b}(t)\}}, \quad (1)$$

where $X_{i,a}(t)$ is the probability that agent i chooses action $a \in \{C, D\}$ at time t . The learning rate, $0 < \alpha_i < 1$, regulates how much new information is used to update the agents' strategies. The expected reward $R_{i,a}(t)$ for agent i choosing action a depends on the opponent's strategy $X_{-i}(t)$. The expected reward for choosing cooperation (C) is

$$R_{i,C}(t) = X_{-i,C}(t) \cdot (1 - g_i) + X_{-i,D}(t) \cdot (-f_i), \quad (2)$$

150 whereas the expected reward for choosing defection (D) is

$$R_{i,D}(t) = X_{-i,C}(t). \quad (3)$$

To address our second question about the importance of agent heterogeneity for explaining the empirical data, we model the agents as either identical (SYM) or heterogeneous (HET).

- **Symmetric agents** (SYM). Agents are symmetric if they share a common learning rate ($\alpha_1 = \alpha_2 \equiv \alpha$), common initial conditions ($X_{1,D}(0) = X_{2,D}(0)$), and common payoff parameters ($f_1 = f_2, g_1 = g_2, r_1 = r_2$).
- **Heterogeneous agents** (HET). Agents are heterogeneous if they differ in learning rates ($\alpha_1 \neq \alpha_2$), initial conditions ($X_{1,D}(0) \neq X_{2,D}(0)$), or payoff parameters ($f_1 \neq f_2, g_1 \neq g_2, r_1 \neq r_2$). As the same cumulative learning dynamics arise when agents swap their parameters, the probabilistic programming procedure cannot converge. To address this identifiability problem, we enforce $\alpha_1 \geq \alpha_2$ and $g_1 \geq g_2$ by sampling α_1 and g_1 from $\text{Unif}[0, 1]$ and multipliers $m_\alpha, m_g \sim \text{Unif}[0, 1]$, then setting $\alpha_2 = \alpha_1 \cdot m_\alpha$ and $g_2 = g_1 \cdot m_g$.

165 Last, we assume that higher defection leads to higher emissions. We define each agent's defection level at time t as $d_i(t) \equiv X_{i,D}(t)$ and the total defection from both agents as $d_{\text{tot}}(t) = d_1(t) + d_2(t)$. To map defection to emissions, we scale behavior by a *carbon-intensity parameter* c , which represents the amount of CO_2 emitted per unit of defection. We assume c is constant over the entire 1850–2022 period; allowing it to vary over time would risk parameter identifiability issues, where this physical scaling could absorb the variance that the behavioral learning dynamics are meant to explain. Real CO_2 emissions are noisy; the parameter $\sigma > 0$ captures this unexplained variation. The observed emissions $E(t)$ are generated via a linear-Gaussian observation model: at each time step, emissions are normally distributed with mean $c d_{\text{tot}}(t)$ and standard deviation σ ,

$$E(t) | \theta \sim \mathcal{N}(c d_{\text{tot}}(t), \sigma^2), \quad t = 1, \dots, T, \quad (4)$$



170 where θ denotes the full set of model parameters and $d_{\text{tot}}(t)$ is computed deterministically from θ via Eq. (1). Eq. (4) links the latent strategies from Eq. (1) to measurable CO₂ emissions. Thus, in total, we test 10 process-based model parameterizations,

$$\{\text{CB}, \text{FGpd}, \text{FGch}, \text{FGsh}, \text{FGha}\} \otimes \{\text{SYM}, \text{HET}\}. \quad (5)$$

2.2 Probabilistic programming

175 *Probabilistic programming* (PP) automates Bayesian inference to make it efficient and accessible. PP enables us to express the foundational equations of our theory as a *data-generating process* (DGP) and then condition that process on empirical observations. In PP, the model runs in two directions: in the *forward* direction, the algorithm draws parameters from prior distributions, propagates them through the model equations, and generates synthetic outcomes. In the *backward* direction, given observed data D , we update uncertainty about parameters θ via Bayes' theorem:

$$p(\theta | D) = \frac{p(D | \theta)p(\theta)}{p(D)}, \quad \text{with} \quad p(D) = \int p(D | \theta)p(\theta) d\theta, \quad (6)$$

where $p(\theta)$ is the prior, $p(D | \theta)$ is the likelihood, and $p(D)$ is the *marginal likelihood* (also called model evidence).

180 The evidence $p(D)$ requires integrating over the full parameter space and is often analytically intractable in high dimensions. Instead of evaluating the posterior in closed form, we use simulation-based methods to sample from it and characterize plausible posterior distributions. We employ the Python package *NumPyro* (Phan et al., 2019), a framework that automates PP using Markov chain Monte Carlo (MCMC) and Hamiltonian Monte Carlo (HMC). HMC uses gradients to move through parameter space, proposing states that respect the geometry of the posterior distribution and improving acceptance rates. Tuning the hyperparameters of HMC (e.g., step size and path length) is challenging; we therefore use the No-U-Turn Sampler (NUTS), which automatically adapts these settings during sampling. We ran 4 independent chains, each with 1,000 warmup steps and 1,000 posterior samples, yielding 4,000 total samples per model. Convergence was verified by ensuring the Gelman–Rubin diagnostic \hat{R} is close to one for all parameters. *NumPyro* also uses *JAX* (Bradbury et al., 2018) to accelerate computations. Model comparison relies on estimating how well each model would predict unseen data. We used leave-one-out cross-validation (LOO-CV), which approximates this by systematically holding out each observation and measuring how accurately the fitted model predicts it (Vehtari et al., 2017). To rank the candidate game structures, we computed Bayesian stacking weights, which assign each model a weight proportional to its predictive performance so that the weights across all models sum to one (Yao et al., 2018). All model comparison diagnostics were computed with *ArviZ*, an open-source Python library for the analysis of Bayesian models (Kumar et al., 2019).

195 The following steps describe how we embed our model into PP.



Table 1. Model parameters and priors

Block	Quantity	Prior / Support	Notes
A. Agents			
<i>A1. Heterogeneous (HET)</i>			
	Learning rates α_i	Unif[0, 1]	Independent per agent
	Initial defection $d_i^{(0)}$	Unif[0, 1]	Independent per agent
<i>A2. Symmetric (SYM)</i>			
	Learning rate α	Unif[0, 1]	Shared across agents
	Initial defection $d^{(0)}$	Unif[0, 1]	Shared across agents
B. Payoffs / regime priors			
	Stag Hunt	$g \sim \text{Unif}[-1, 0], f \sim \text{Unif}[0, 1]$	Enforces $g < 0, f > 0$
	Chicken	$g \sim \text{Unif}[0, 1], f \sim \text{Unif}[-1, 0]$	Enforces $g > 0, f < 0$
	Prisoner’s Dilemma	$g \sim \text{Unif}[0, 1], f \sim \text{Unif}[0, 1]$	Enforces $g > 0, f > 0$
	Harmony	$g \sim \text{Unif}[-1, 0], f \sim \text{Unif}[-1, 0]$	Enforces $g < 0, f < 0$
C. Common parameters			
	Emission scale c	Unif[3, 30]	Shared across agents
	Observation noise σ^a	Normal(1×10^{10})	Normal prior on noise scale

^a The model is fit to absolute metric tons of CO₂ (scaled internally by 10¹⁰). The Normal scale parameter 1×10^{10} metric tons (1 GtCO₂ in scaled units) places most prior mass below ~ 6 GtCO₂, providing a weakly informative prior that comfortably covers plausible observation errors given recent global emissions of approximately 40 GtCO₂/yr.

200

205

- 1. Data generating process (DGP):** Mathematical formulations that specify how inputs generate outputs under our theory. Our *DGP* includes: the fear-greed payoff definition (Eqs. (2)–(3)); agent behavior via Eq. (1); the mode (SYM or HET), which sets α_i and initial strategies; and the conversion of defection levels to CO₂ emissions via Eq. (4).
- 2. Priors:** In general, priors (Table 1) represent our knowledge about the parameters; however, to remain agnostic about the possible scenarios, we consider uniform distributions for all parameters except the observation noise. We place priors over g, f, α_i, c and σ to encode initial uncertainty. For σ , we use a weakly informative Normal(1×10^{10}) prior (in metric tons), which concentrates mass on plausible noise levels while excluding negative values. Given that recent annual global emissions are approximately 40 GtCO₂, this prior comfortably covers any plausible observation error.
- 3. Data:** We use annual global CO₂ emissions, $E(t)$, from fossil fuel and land use change in the period 1850–2022 ($N = 173$ observations) (Friedlingstein et al., 2022).



4. **Likelihood:** Assuming conditionally independent observations as specified in Eq. (4), the full time-series likelihood is the product of the individual likelihoods:

$$p(E_{1:T} | \theta) = \prod_{t=1}^T \mathcal{N}(E(t); cd_{\text{tot}}(t), \sigma^2), \quad (7)$$

where θ represents all model parameters and $d_{\text{tot}}(t)$ is computed deterministically from θ via Eq. (1).

210 5. **Posterior inference:** We infer the posterior distribution over the model parameters, $p(\theta | E_{1:T}) \propto p(E_{1:T} | \theta)p(\theta)$, using the No-U-Turn Sampler (NUTS), a variant of Hamiltonian Monte Carlo, implemented in NumPyro. This distribution represents the plausible values for the parameters, given the observed CO₂ emissions. We verify the reliability of the inference using standard convergence diagnostics, confirming that independent chains have converged ($\hat{R} \approx 1$), the effective sample size is large, there are no NUTS divergences, and autocorrelation between samples is low.

215 6. **Model comparison:** We compare models using Pareto-smoothed importance sampling leave-one-out cross-validation (PSIS-LOO) (Vehtari et al., 2017). This method has two key metrics: the Expected Log Pointwise Predictive Density (ELPD), where higher values indicate better predictive accuracy, and the effective number of parameters (p_{loo}), which estimates model complexity. To rank models, we compute *Bayesian stacking weights* (Yao et al., 2018), which find the optimal model combination by maximizing leave-one-out predictive performance. The weights sum to 1 and indicate
220 each model's contribution to out-of-sample prediction. Additionally, we evaluate point-prediction accuracy using Root Mean Square Error (RMSE):

$$\text{RMSE} = \sqrt{\frac{1}{T} \sum_{t=1}^T (E_{\text{obs}}(t) - E_{\text{model}}(t))^2}. \quad (8)$$

2.3 Parameter sensitivity analysis

We conducted a sensitivity analysis using a posterior-preserving perturbation approach to identify which parameters most
225 strongly influence emission projections. For each model, posterior mean parameters were individually perturbed by $\pm 20\%$, and the resulting change in cumulative 2023–2100 emissions was measured. This approach preserves the structure of the joint posterior while testing the standalone impact of individual variables.

The sensitivity index S_θ for a parameter θ is computed as the maximum relative change in cumulative emissions from the baseline:

$$230 S_\theta = \max_{k \in \{0.8, 1.2\}} \left(\frac{|E_{k\theta} - E_{\text{baseline}}|}{E_{\text{baseline}}} \right), \quad (9)$$

where $E_{k\theta}$ denotes cumulative emissions under perturbation factor k (representing a $\pm 20\%$ shift). An S_θ of 20% indicates a proportional (linear) response to the parameter change, as observed for carbon intensity. Values exceeding 20% indicate that the parameter amplifies emission outcomes.



3 Results

235 We evaluated 10 model configurations, including two parameterizations (Cost-Benefit and Fear-Greed) and two agent structures (symmetric and heterogeneous). Table 2 summarizes predictive performance via PSIS-LOO cross-validation for the eight non-Chicken configurations. The fitting procedure fails to converge for the Chicken game because of bimodality: two feasible solutions exist in the parameter space. Its results are reported separately in Sect. 3.3, since the ELPD scores are inflated by highly informed priors. Three patterns emerge:

- 240
1. **Prisoner’s Dilemma, Cost-Benefit, and Chicken game models explain emissions best.** They substantially outperform coordination games (Stag Hunt, Harmony). The top two models are variants of the Prisoner’s Dilemma and the Chicken game; however, the Chicken game requires clustering, and its apparent better fit is inflated by informative priors.
 2. **Heterogeneity consistently improves predictive accuracy.** For every game, allowing agents to differ in their strategic parameters yields large Expected Log Predictive Density (ELPD) gains (+6.85 to +122.0), with the exception of
245 Harmony (−0.30). The heterogeneous PD dominates, receiving 99.4% of the Bayesian stacking weight.
 3. **Sensitivity is concentrated in the dominant emitter.** In the top model, the learning rate and greed parameter of Agent 1 (79% of emissions) are the strongest leverage points. A 20% perturbation in Agent 1’s learning rate changes cumulative future emissions by up to 32.0% (see Fig. 4).

The rest of this section unpacks these findings. We first describe the dominant heterogeneous Prisoner’s Dilemma and
250 formally rule out the coordination games. We then evaluate the simpler Cost-Benefit framework to show why a single-parameter model is sufficient for fitting historical emissions but lacks diagnostic utility for future interventions. Next, we examine the Chicken game as the only structurally appropriate alternative and resolve its bimodal posterior into two distinct behavioral regimes. Finally, we map the systemic leverage points through a cross-game sensitivity analysis to identify which parameters control the emission path regardless of the assumed game. All reported values are posterior medians, with 95% Credible
255 Intervals (CI) provided where relevant.



Table 2. Bayesian model comparison ranked by ELPD (Chicken heterogeneous excluded; see Sect. 3.3). Δ ELPD is the difference from the top model. Wt. = Bayesian stacking weight over these eight models; the heterogeneous Chicken model is excluded because its bimodal posterior implies regime-conditioned priors that are not directly comparable to the uniform priors used here. p_{100}/k = effective/nominal parameter ratio; values < 0.5 indicate structural misspecification. Weights may not sum exactly to 1.00 due to rounding; the residual is distributed across lower-ranked models in amounts < 0.01 .

Rank	Model	ELPD	Δ ELPD	Wt.	$\frac{p_{100}}{k}$	Interpretation
<i>Tier 1: Top performers</i>						
1	PD (Heterogeneous)	158.5	0.0	0.99	0.68	Best fit; dominant weight.
2	Cost-Benefit (Het.)	155.9	-2.6	0.00	0.74	Redundant; duplicates PD predictions.
<i>Tier 2: Lower accuracy</i>						
3	PD (Symmetric)	151.6	-6.8	0.00	0.83	Underfit; symmetric constraint too strong.
4	Cost-Benefit (Sym.)	150.7	-7.8	0.01	0.86	Complementary; captures unique variance.
<i>Tier 3: Structural failure</i>						
5	Stag Hunt (Het.)	76.6	-81.9	0.00	0.24	Bad fit; data contradict coordination.
6	Stag Hunt (Sym.)	-45.4	-203.9	0.00	0.50	Worse than intercept-only baseline.
7	Harmony (Sym.)	-268.2	-426.7	0.00	0.30	Structurally incompatible.
8	Harmony (Het.)	-268.5	-427.0	0.00	0.28	Structurally incompatible.

3.1 The primary diagnosis: heterogeneous Prisoner’s Dilemma and excluded models

The top-ranked model identifies a heterogeneous Prisoner’s Dilemma, where both agents exhibit positive fear and greed but possess fundamentally distinct strategic profiles. Agent 1 (accounting for 79.2% of emissions) is greed-driven ($g_1 = 0.34$, 95% CI [0.07, 0.87]) and adapts rapidly ($\alpha_1 = 0.27$, 95% CI [0.08, 0.78]), defecting primarily due to the temptation to exploit. Agent 2 (20.8% of emissions) is fear-driven ($f_2 = 0.64$, 95% CI [0.17, 1.00]) and adapts slowly ($\alpha_2 = 0.13$, 95% CI [0.00, 0.36]), defecting defensively. This asymmetry, which symmetric models fail to capture, drives the superior predictive accuracy of heterogeneous configurations (Table C3). The emergent 79%/21% emission split closely matches historical estimates of the Global North–South divide (Hickel, 2020; Friedlingstein et al., 2022).

In contrast, coordination games (Stag Hunt, Harmony) are structurally incompatible with the data. Symmetric versions yield negative ELPD values, and heterogeneous versions severely underperform the PD (Δ ELPD ≤ -82.5). Full posterior estimates and convergence diagnostics for all models are provided in Appendix A.

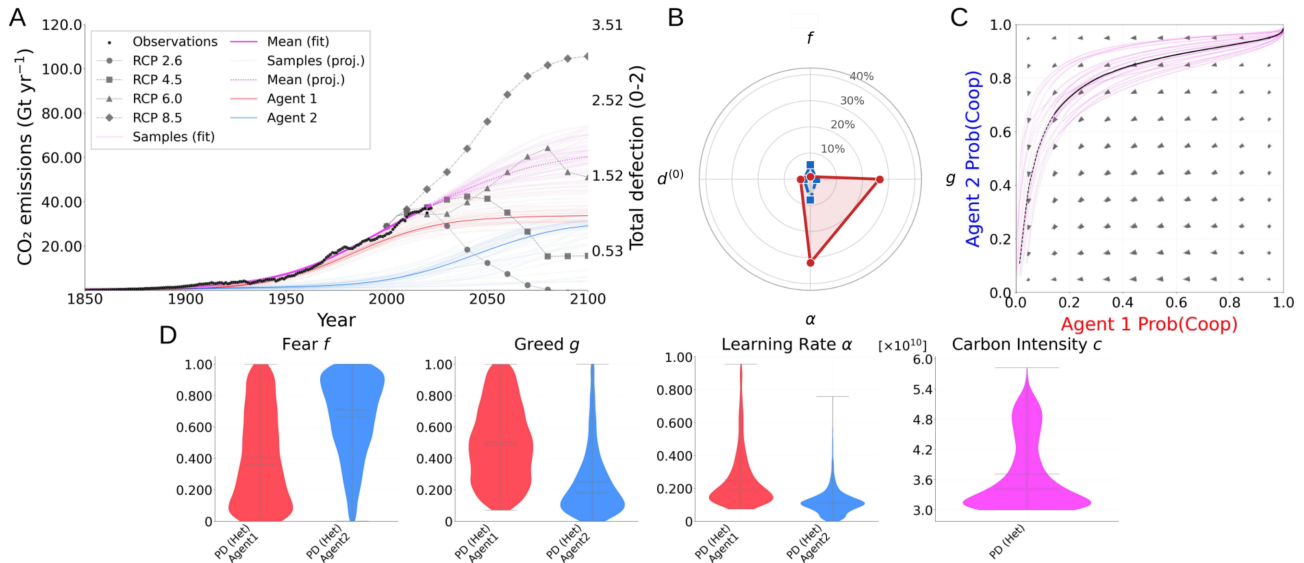


Figure 2. Dynamics of the heterogeneous Prisoner's Dilemma (99.4% Bayesian stacking weight). **(a)** Historical fit and future projections. Black dots: observed global CO₂ emissions (1850–2022). Magenta traces: posterior predictive samples; solid magenta: posterior median fit. Dotted magenta: median projection under status-quo parameters. Red and blue solid lines: individual emission trajectories of Agent 1 (dominant emitter, 79% of cumulative emissions) and Agent 2 (smaller emitter, 21%), respectively. Right-axis labels indicate total defection levels (d_{tot}). Dashed grey lines: Representative Concentration Pathways (RCP 2.6, 4.5, 6.0, 8.5); the model projection aligns most closely with RCP 6.0, corresponding to approximately 2.0–3.7°C warming by 2100 (best estimate $\sim 2.8^\circ\text{C}$; IPCC AR5). **(b)** Sensitivity analysis (radar plot). Each axis represents the percentage change in cumulative 2023–2100 emissions resulting from a $\pm 20\%$ perturbation of the corresponding parameter at its posterior median. Red dots: Agent 1 parameters; blue dots: Agent 2 parameters. Agent 1's learning rate (α_1 , 32.0%) and greed (g_1 , 26.4%) are the two strongest policy leverage points, both exceeding the proportional benchmark of 20%. Agent 2's learning rate (α_2 , 7.9%) has moderate but secondary sensitivity. **(c)** Phase-space dynamics. Each axis represents an agent's probability of cooperation. Grey arrows: deterministic flow field under posterior median parameters. Magenta traces: 50 stochastic posterior trajectories. The solid black line: median trajectory. The dashed lines represent the projections. Nearly all trajectories converge toward the lower-left corner (mutual defection), confirming that the Prisoner's Dilemma equilibrium is the sole attractor of the learning dynamics. **(d)** Posterior distributions of key parameters. Violin plots show the marginal posteriors for fear (f), greed (g), learning rate (α), and carbon intensity (c) for Agent 1 (red) and Agent 2 (blue). Agent 1 exhibits higher greed ($g_1 = 0.34$) and a faster learning rate ($\alpha_1 = 0.27$), consistent with a temptation-driven, rapidly adapting emitter. Agent 2 exhibits higher fear ($f_2 = 0.64$) and a slower learning rate ($\alpha_2 = 0.13$), consistent with a vulnerability-driven, structurally constrained emitter. The carbon intensity posterior (magenta) concentrates near the lower bound ($c \approx 3.4$).

3.2 The parsimonious alternative: Cost-Benefit framework

While the heterogeneous Prisoner's Dilemma (PD) represents the best fit, the Cost-Benefit (CB) parameterization reveals a useful structural insight. The CB framework is a constrained version of the PD in which the fear of exploitation and the greed to free-ride collapse into a single cost-benefit ratio ($f = g = r = c/b$). This parameterization achieves the second-highest



predictive accuracy overall (ELPD 155.9), with only a $\Delta\text{ELPD} = -2.6$ deficiency. Its Bayesian stacking weight of 0.00 reflects predictive redundancy rather than poor performance: it generates nearly identical trajectories to the top-ranked model (Fig. B1a). The constrained symmetric CB model has a non-zero stacking weight (0.01) alongside an exceptional p_{100}/k of 0.86, indicating that it captures unique dynamics (Fig. B1b).

275 The inferred parameters of the heterogeneous CB model mirror the asymmetric structure of the top-ranked PD. Agent 1 has a higher cost-benefit ratio ($r_1 = 0.29$, 95% CI [0.08, 0.86]) and adapts rapidly ($\alpha_1 = 0.29$, 95% CI [0.08, 0.85]), accounting for 90.8% of cumulative emissions. Agent 2 has a lower cost-benefit ratio ($r_2 = 0.13$, 95% CI [0.00, 0.56]) and a much slower learning rate ($\alpha_2 = 0.12$, 95% CI [0.00, 0.53]), contributing only 9.2% of emissions. The carbon intensity c : all top-performing models has median values in the range 3.0–5.1, whereas structurally incompatible models inflate c to ≈ 10 –12 with wide
280 posteriors (95% CI up to 27), indicating that c is absorbing misfit rather than reflecting genuine carbon intensity.

The statistical redundancy in the Cost-Benefit variant highlights the distinction between predictive sufficiency and diagnostic utility. The Cost-Benefit model is statistically sufficient to fit historical emissions. However, when evaluating future policy interventions, this single-parameter formulation is inadequate. By merging fear and greed into a single ratio r , the CB framework combines two entirely different policy levers: sensitivity analysis confirms that r_1 shifts cumulative emissions by 23.4%, but
285 this effect of temptation and fear cannot be decomposed (Table D1). The unconstrained PD is therefore superior not just for its marginally better fit, but because it distinguishes between the underlying strategic incentives.



3.3 Chicken game bimodality

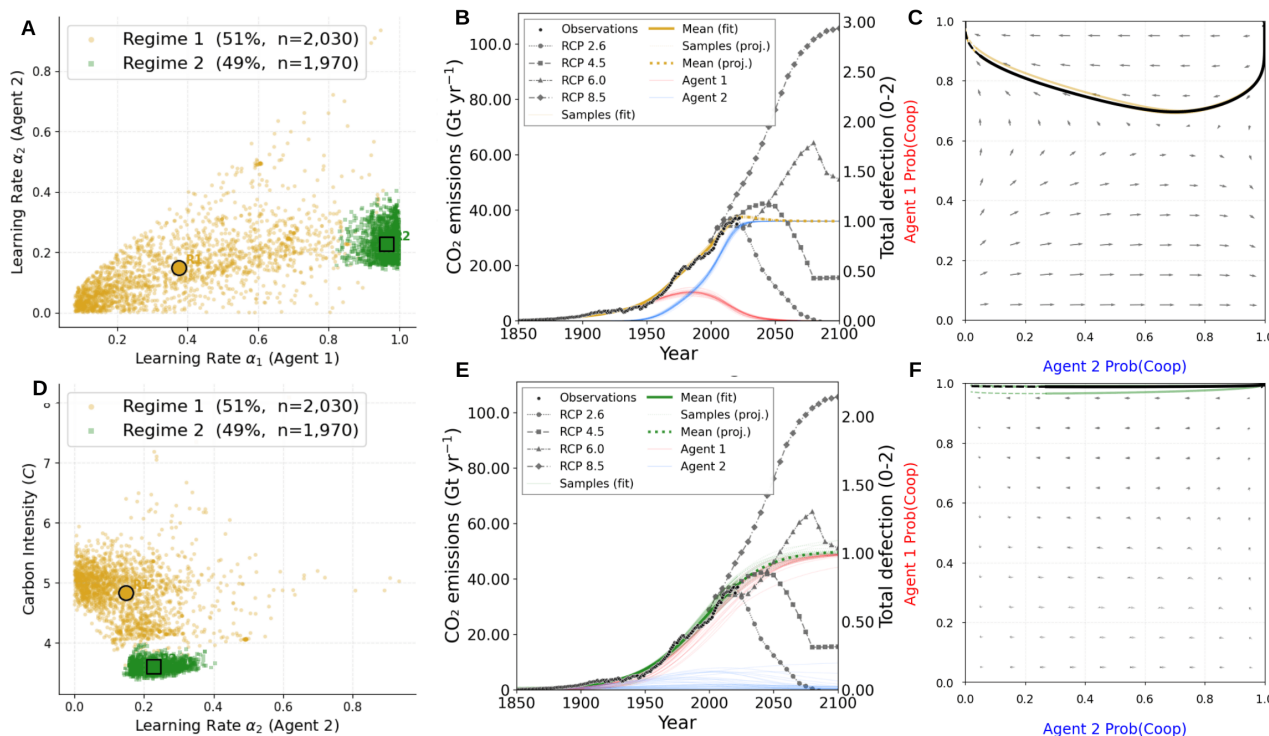


Figure 3. Bimodality in the heterogeneous Chicken game. (a) & (d) Posterior parameter distributions revealing two distinct, behavioral regimes with equal posterior share. Panel (a) shows the clustering of learning rates (α_1 vs. α_2); Panel (d) maps carbon intensity (c) against α_1 . Regime 1 (yellow): a fast-adapting, greedy Agent 1 in a low-carbon economy. Regime 2 (green): a moderate-adapting Agent 1 in a high-carbon economy. (b) & (e) Historical fits and future projections for Regime 1 (b) and Regime 2 (e). In both regimes the dominant emitter saturates at near-full defection while the smaller agent swerves toward cooperation. Regime 1 plateaus near 37 Gt CO_2/yr ; Regime 2 grows to about 50 Gt CO_2/yr . (c) & (f) Phase-space dynamics showing the strategy flows for Regime 1 (c) and Regime 2 (f), reflecting the distinct trajectories.

The heterogeneous Chicken model does not converge as a single posterior (\hat{R} up to 1.83, ESS < 60; Table G2). Bayesian Gaussian mixture clustering on α_1 , α_2 , and c splits the posterior into two regimes with equal mass (silhouette score 0.804; Table 3). We exclude both from Table 2 because their highly informed priors inflate the ELPD.

Both regimes reach the same qualitative outcome: the dominant emitter ends at full defection and the smaller agent swerves toward cooperation. They differ in how they trade off behavioral speed against physical scaling. Regime 1 pairs a fast-adapting, greedy Agent 1 ($\alpha_1 = 0.97$, 95% CI [0.89, 1.00]; $g_1 = 0.55$, 95% CI [0.46, 0.63]) with a low-carbon economy ($c = 3.58$, 95% CI [3.45, 3.74]). This gives a 54%/46% emission split and a projected plateau near 37 Gt CO_2/yr . Regime 2 pairs a moderately adapting Agent 1 ($\alpha_1 = 0.41$, 95% CI [0.13, 0.60]; $g_1 = 0.38$, 95% CI [0.15, 0.69]) with a high-carbon economy ($c = 4.87$,



95% CI [4.50, 5.50]), giving a 75%/25% split and a higher plateau near 50 Gt CO₂/yr (Fig. 3; Table 3). The same aggregate trajectory admits either factorization.

300 Regime 1’s better fit comes from three factors. Its fast adaptation rate ($\alpha_1 \approx 1$) rigidly tracks every historical fluctuation. Negative fear lets Agent 2 pivot easily once Agent 1 reaches full defection, which simplifies the projections. The low scaling factor (c) artificially shrinks the inferred noise (σ). Combined with the informed priors, these effects tighten the historical fit.

Regime 2 behaves much like the heterogeneous Prisoner’s Dilemma (PD) model. Their core metrics match closely: adaptation speeds ($\alpha_1 = 0.41$ vs. 0.27; $\alpha_2 = 0.11$ vs. 0.13), dominant greed ($g_1 = 0.38$ vs. 0.34), emission shares (75/25 vs. 79/21), and ELPD scores (159.3 vs. 158.5). The only structural difference is the sign of the fear parameter, and fear has little weight in either model ($f_1 < 8\%$).

305 The comparison between these regimes and the other models is not on equal terms. PD-Het converges under uniform priors, while Regime 1 only emerges after clustering a non-converged posterior to inform its priors. Even so, three independent specifications (the heterogeneous PD, Regime 2, and the heterogeneous Cost-Benefit model) recover the same asymmetric, greed-driven structure. That structurally different models reproduce the same result suggests the finding is driven by the data, not by the choice of model.

Table 3. Chicken game regimes: posterior summaries, emission shares, and projection scale. Both regimes converge to the same asymmetric Chicken equilibrium (Agent 1 saturates at full defection, Agent 2 swerves toward cooperation), but at different physical and behavioral scales.

	Regime 1	Regime 2
<i>Behavioral (median and 95% CI)</i>		
α_1	0.97 [0.89, 1.00]	0.41 [0.13, 0.60]
α_2	0.23 [0.15, 0.36]	0.11 [0.00, 0.31]
g_1	0.55 [0.46, 0.63]	0.38 [0.15, 0.69]
g_2	0.28 [0.17, 0.39]	0.31 [0.10, 0.70]
f_1	-0.97 [-1.00, -0.89]	-0.75 [-1.00, -0.09]
f_2	-0.73 [-1.00, -0.44]	-0.48 [-1.00, -0.07]
<i>Physical (median and 95% CI)</i>		
c	3.58 [3.45, 3.74]	4.87 [4.50, 5.50]
σ	0.074	0.094
<i>Emission structure</i>		
Agent 1 share	54%	75%
Agent 2 share	46%	25%
<i>Fit quality (1850–2022)</i>		
ELPD	199.2	159.3
RMSE (Gt CO ₂ /yr)	0.73	0.92
Projected plateau	~37 Gt CO ₂ /yr	~50 Gt CO ₂ /yr



310 3.4 Systemic leverage points: cross-game sensitivity analysis

Although the Prisoner's Dilemma dominates the stacking weights, the sensitivity structure across the heterogeneous models reveals which parameters would matter *if* a different game governed the climate problem (Fig. 4). Three patterns appear. First, in the PD, CB, and Stag Hunt models, Agent 1's learning rate (α_1) is the most sensitive behavioral parameter.

Depending on the game, the *strategic* parameters would be different. In the Prisoner's Dilemma, greed dominates fear: g_1 shifts cumulative emissions by 26.4%, whereas f_1 produces only 1.0%. The system is controlled by the dominant emitter's temptation to exploit (Dorninger et al., 2021).

The Stag Hunt presents a fundamentally different picture. Greed is almost zero ($g_1 = 0.2\%$, $g_2 = 0.1\%$), meaning agents are not tempted to exploit ($g < 0$). Instead, the system is governed by Agent 1's fear ($f_1 = 21.9\%$), Agent 2's initial conditions ($d_2^{(0)} = 24.9\%$), and both agents' fear parameters ($f_2 = 8.4\%$). The dominance of initial conditions means that in a coordination game, the barrier to cooperation is not temptation but path dependency. In the Stag Hunt game, the dominant emitter's fear sensitivity ($f_1 = 21.9\%$) would correspond to a world where lack of trust, not temptation, is the central problem. However, the data reject this diagnosis: the Stag Hunt's RMSE (1.53×10^9 ; Eq. (8)) is significantly higher than the PD's (9.29×10^8), and its posterior has extreme parameter values ($f_1 \approx 0.99$, $\alpha_1 \approx 0.99$), indicating that the model fits the emissions only by relying on extreme fear.

Similar to PD, in both Chicken regimes greed dominates. A 20% perturbation in g_1 shifts cumulative emissions by 23.1% in Regime 1 and 32.2% in Regime 2 which is comparable to the PD. The fear parameters have a low sensitivity ($f_1 < 8\%$ in both regimes). Figure 4 shows the similarity between the Chicken and Prisoner's Dilemma sensitivity profiles, contrasted with Stag Hunt's profile. The bimodality of the heterogeneous Chicken model is examined in Sect. 3.3.

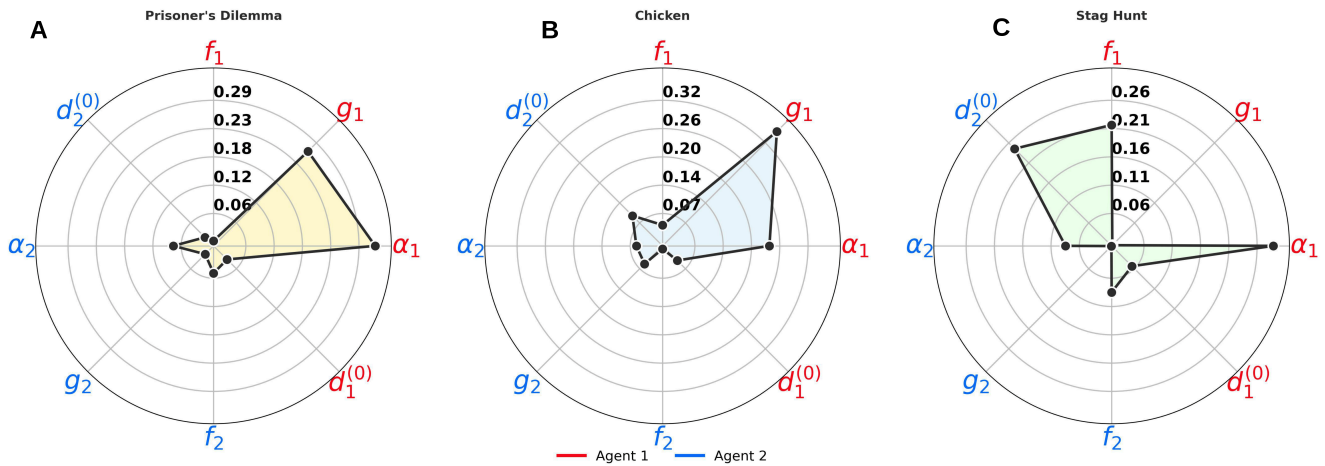


Figure 4. Cross-game sensitivity analysis for heterogeneous models. Each radar plot shows the percentage change in cumulative 2023–2100 emissions from a $\pm 20\%$ perturbation of each parameter at its posterior median. Red labels indicate Agent 1 parameters; blue labels indicate Agent 2 parameters. **Panel (a) (Prisoner’s Dilemma):** Greed (g_1) and learning rate (α_1) dominate; fear (f_1) is secondary. **Panel (b) (Chicken, Regime 2):** Agent 1’s greed (g_1) dominates, with the slower agent’s learning rate (α_2) as the secondary lever. Agent 1’s learning rate (α_1) collapses to near zero because Agent 1 reaches a corner solution ($d_1 \approx 1$) early in the projection horizon. Fear parameters remain small. Regime 1 shows the same greed-dominant pattern (Table D3). **Panel (c) (Stag Hunt):** Fear (f_1), initial conditions ($d_2^{(0)}$), and learning rate (α_1) dominate; greed is structurally absent ($g < 0$). Across all three panels, Agent 2’s learning rate (α_2) exhibits at least moderate sensitivity.

4 Discussion

330 Our results suggest a structural answer to the persistence of the climate cooperation problem. The data favor a free-riding problem among fundamentally different actors with different shares in historical emissions and different responsibilities.

4.1 Historical CO₂ emission problem is a free-riding problem, not a coordination failure

Our results indicate that coordination games and Harmony cannot reproduce the observed trajectory of climate negotiations. Instead, the heterogeneous Prisoner’s Dilemma shows that both parties have incentives favoring defection over cooperation.

335 Without penalties, free-riding has the best economic gains, so defection becomes rational.

Therefore, self-enforcing environmental agreements face a structural paradox. Exactly when deep emission cuts are needed most, the strong incentive to free-ride ensures that voluntary treaties can only sustain minimal commitments (Barrett, 2003). Because of this persistent free-rider dynamic, framing sustainability problems as a simple coordination game may not only be too optimistic but potentially damaging, since short-term interests will inevitably introduce PD dynamics even in those settings

340 (Carrozzo Magli and Manfredi, 2022).



Our results also diverge from the institutional perspective, which shows that communities can overcome commons dilemmas through self-governance without external enforcement (Ostrom, 1990). Our inferred parameters provide direct empirical evidence against this possibility in the climate case: the posterior concentrates on positive greed and positive fear, precisely the conditions under which the dominant strategy is defection regardless of the other’s behavior. Whatever self-governing mechanisms have operated over the past years, they have not been sufficient to shift the effective incentive structure away from a Prisoner’s Dilemma.

Our findings highlight a global free-rider problem, which may appear to conflict with theories emphasizing domestic politics. Aklin and Mildenerger (2020) argue that in important scenarios free-riding is “empirically unsubstantiated” and that climate politics are driven entirely by domestic distributive conflict. Their case-study evidence suggests that countries base their climate policies on internal dynamics, not on the actions of others. However, our analysis operates at a different level: even if domestic politics are the proximate cause of individual policy choices, the aggregate global outcome of those choices is best explained by Prisoner’s Dilemma dynamics. The two perspectives are therefore complementary rather than contradictory. Domestic distributive conflict may be the mechanism through which free-riding incentives are expressed in practice. Even if states do not consciously intend to free-ride, the overarching incentive structure they collectively face produces the same trajectory. The PD framing captures this emergent strategic structure, regardless of the micro-level mechanisms driving it.

The need to target free-rider incentives remains consistent across plausible game structures. Even under the Chicken game, the most optimistic alternative in which agents prefer unilateral cooperation over mutual defection, greed reduction remains the dominant policy lever, while fear-based interventions have a negligible effect (Fig. 4; Table D3). This finding aligns with classic arguments that climate change is fundamentally a collective action problem (Levin et al., 2012) and supports the broader theoretical literature on global public goods assuming free-riding incentives (Buchholz and Sandler, 2021).

4.2 Asymmetric incentives and the greed-fear decomposition

Models with asymmetric agents clearly outperform those with identical agents: the heterogeneous Prisoner’s Dilemma gains +6.45 ELPD over its symmetric counterpart, and across all game types, heterogeneity consistently improves predictive accuracy (Table C3). Although the model infers two abstract agents from aggregate data alone and does not observe the North-South divide directly, the inferred profiles reveal a clear structural interpretation. The dominant emitter, contributing 79% of emissions, is driven by the temptation to free-ride on others’ cooperation, consistent with the historical pattern of industrialized nations profiting from non-cooperation (Dorninger et al., 2021). The smaller emitter defects out of fear, driven by vulnerability to exploitation and limited capacity to adapt. Their learning speeds differ accordingly: the dominant emitter responds quickly to economic incentives, while the smaller emitter is structurally slower to shift its strategic posture.

The inferred 79%/21% emission split aligns closely with historical accounting of the North-South divide. Hickel (2020) found that the Global North was responsible for 92% of CO₂ emissions in excess of the planetary boundary (350 ppm) between 1850 and 2015, and that Annex I nations accounted for 90% of these excess emissions. The Annex I share of cumulative historical fossil CO₂ from the Global Carbon Project dataset (Friedlingstein et al., 2022) is approximately 74%, compared to Agent 1’s inferred 79%. This comparison is suggestive and should be interpreted cautiously: the model was not provided



375 with any disaggregated emissions data, and the correspondence does not mean a validation of the North-South mapping but it
motivates the interpretation developed here. The mapping onto the Global North and South is our interpretation, motivated by
the close match between the model's inferred emission shares and the historical record.

The fear-greed decomposition, introduced by Coombs (1973) and widely used in learning-dynamics analyses of social
dilemmas (Macy and Flache, 2002), helps explain why the Prisoner's Dilemma is uniquely difficult to resolve: it activates
380 *both* motives for defection at once. In a Stag Hunt, actors defect only because they fear others will not cooperate; in a Chicken
game, they defect only because they are tempted to exploit others' cooperation. In the PD, both motives operate simultaneously,
which means that addressing one alone may not be sufficient. Experimental evidence supports this distinction: Ahn et al. (2001)
showed that fear and greed are both significant predictors of cooperation rates in PD games. Our sensitivity analysis confirms
this pattern at the global scale (Fig. 4; Appendix D): a 20% perturbation in Agent 1's greed (g_1) shifts cumulative emissions
385 by 26.4%, whereas the same perturbation in Agent 1's fear (f_1) produces only 1.0%.

These asymmetric profiles provide a game-theoretic grounding for the principle of Common but Differentiated Responsibil-
ities (United Nations, 1992). The greed-driven agent resembles the Global North: historically the largest cumulative emitter,
benefiting economically from fossil fuel use, and defecting because it profits from doing so. The fear-driven agent resembles
the Global South: bearing a smaller share of historical emissions, facing greater climate vulnerability, and defecting not out of
390 greed but because it defends itself from complete exploitation. Our two-agent framework captures this North-South asymmetry
in stylized form, but the reality is more complex.

4.3 Projection: status-quo alignment with RCP 6.0

The social dilemma models align most closely with RCP 6.0 (Appendix E), which corresponds to approximately 2.0–3.7°C
warming by 2100 (best estimate $\sim 2.8^\circ\text{C}$; IPCC AR5). The only exception is the historically implausible Chicken Regime 1,
395 which aligns with RCP 4.5 (Sect. 3.3). This warming range is broadly consistent with independent assessments from a different
scenario family. For instance, the UNEP Emissions Gap Report projects approximately 3.1°C under current policies (United
Nations Environment Programme, 2024), and the IPCC AR6 SSP2-4.5 scenario provides a best estimate of 2.7°C (range
2.1–3.5°C).

Our game-theoretic models calibrated purely on historical emissions arrive at the same warming range as physics-based
400 integrated assessment models. This agreement strengthens both the strategic interpretation and the temperature projection. At
this level of warming, Armstrong McKay et al. (2022) identified multiple climate tipping points that become likely. Governing
such tipping elements requires moving beyond pure economic optimization toward sustainability and safe operating space
paradigms, because no policy criterion alone guarantees satisfying the others (Barfuss et al., 2018).

4.4 Prioritizing policy interventions

405 Sensitivity analysis identifies the learning rate and the greed of the dominant emitter as two influential levers of the system.
Our results provide quantitative support: g_1 is the second most sensitive parameter (26.4%), suggesting that even a modest
reduction in the free-riding incentive produces a significant reduction in greenhouse gas emissions. In contrast, the speed of



adaptation by the smaller emitter has a substantially smaller effect than the dominant emitter's ($\alpha_2 = 7.9\%$ vs. $\alpha_1 = 32.0\%$), making it a secondary rather than primary lever. This provides a quantitative justification for focusing on the actions of large emitters.

According to the top performing models to deal with emissions we need a policy sequence. The first priority is to reduce the net benefits of defection (reducing g), which can be approached from two sides: reducing costs of sustainable behavior and establishing enforceable mechanisms. On the incentive side, when clean energy becomes cheaper than continued reliance on fossil fuels, the greed parameter decreases which potentially pushes the game towards a coordination problem or even a Harmony game. The enforcement dimension shows why the voluntary commitments of the Paris Agreement (United Nations, 2015) were not successful since they failed to change the structure of strategic payoffs for members. Binding mechanisms can make the free-riding more costly. For instance, trade penalties for non participating countries (free riders) can help with creation of coalitions, and carbon border adjustment mechanisms can reduce carbon leakage and encourage cooperating countries (Böhringer et al., 2016; Nordhaus, 2015).

Our sensitivity analysis predicts the dynamic outcome of this policy sequence. Suppressing the free-riding incentive shifts the strategic landscape toward a Stag Hunt, where the system's sensitivity to fear is mathematically enhanced (with f_1 jumping to 21.9%). When the temptation is removed, vulnerability becomes the main obstacle to cooperation. After controlling for greed, addressing defection driven by fear through strategic transfers becomes the next step. Since the smaller emitter's fear remains high ($f_2 = 0.64$), it will continue to defect defensively without support. Therefore, Loss and Damage financing, technology transfers, and grants are not charity, but strategic payments necessary to encourage cooperative behavior (Schinko et al., 2019). This is in line with Okada (2023), who showed that combining penalties with transfers enables full cooperation in dynamic climate games. However, our sensitivity ranking warns that reducing fear yields less immediate emission reduction than reducing greed.

4.5 Limitations and extensions

Our framework has some abstractions that limit its current scope but also define a clear agenda for future work. Any model is a simplification and the relevant question for each abstraction is not whether it could be made more realistic, but what specific research question would justify relaxing it. A two-agent model with constant parameters and uniform priors reproduces historical CO₂ emissions and arrives at projections consistent with RCP 6.0, the same warming range that far more complex Earth system models produce. This suggests that even under extensive simplification, the model can still contribute to our understanding of the strategic underpinnings of CO₂ emissions. For each modeling choice below, we explain its purpose in the current framework and identify the question that would motivate an extension.

Two-agent aggregation. Modeling the global landscape using two agents cannot capture coalition formation, bloc fragmentation, or the emergence of major emitters such as China, whose strategic role has changed over time. We chose this parsimony because it makes the role of macro-level inequality tractable without overfitting. We did not define the agent roles as the Global North and South; inequality emerged from the aggregated data. The similarity to the historical North-South divide motivates



our interpretation but should not be taken as a validation. Extending to asymmetric agents with N players enables the study of coalitions at the cost of expanding the parameter space, which requires richer identification strategies.

Constant parameters. All model parameters, including carbon intensity, are held constant over the time. This assumption ignores major institutional, technological, and economic transitions, but serves a specific identification purpose: allowing carbon intensity or game payoffs to change over time runs the risk of absorbing the variance of learning dynamics. Therefore, the inferred values should therefore be interpreted as effective aggregates over time.

Myopic agents and missing feedback. Learning agents are embedded in a single environmental state. This ignores temporal discounting, which is central to sustainability challenges. More importantly, the model lacks a connection between strategic dynamics and the biophysical state: greenhouse gas emissions change the climate, which in turn changes the payoffs. Placing these agents in a co-evolutionary framework with climate tipping elements (Tilman et al., 2020; Weitz et al., 2016; Barfuss et al., 2020) and introducing feedback would bring the model closer to Anthropocene reality.

Observation model. Equation (4) connects the strategic dynamics to physical CO_2 emissions through c , while σ absorbs unexplained variation. We assume that greenhouse gas emissions scale linearly with defection in the mean $cd_{\text{tot}}(t)$, and that deviations from this trend are independent and identically distributed in the product likelihood in Eq. (7). In practice, the relationship between strategic behavior and physical emissions is likely to be nonlinear, and prediction errors may be correlated. Relaxing these assumptions is a natural direction for future work.

Methodological scope. While our approach shows that Bayesian probabilistic programming can bridge stylized game-theoretic models and empirical data, the current application is limited to a single observable (aggregate CO_2 emissions). Incorporating additional data such as disaggregated emissions and physical tipping elements would strengthen parameter identification.

These strategic abstractions yielded a tractable methodological framework, bridging dynamic game-theoretic models with empirical data without introducing excessive structural complexity.

5 Conclusions

By linking stylized social dilemma games to empirical data via Bayesian probabilistic programming, we move from metaphor to an empirically grounded diagnosis (Donges and Barfuss, 2017). The data resolve the contested strategic architecture of the climate problem (Aklin and Mildenerger, 2020): the results suggest that the world is locked in a structurally asymmetric Prisoner's Dilemma, not a coordination failure. A dominant bloc defects out of greed, exploiting the economic benefits of a carbon-intensive free-ride, while a smaller bloc defects defensively out of fear of exploitation. This emergent asymmetry, recovered from aggregate data, provides a quantitative foundation for the principle of Common but Differentiated Responsibilities (United Nations, 1992).

Methodologically, our approach shows that Bayesian probabilistic programming can serve as a bridge between formal narratives –stylized, interpretable models – and empirical data. By treating game-theoretic learning dynamics as a data-generating



process and conditioning on observed emissions, this framework provides posterior distributions for the parameters and enables rigorous model comparison.

475 Based on the formal social dilemma narratives tested, our results suggest a sequenced policy response. The priority is to dismantle the free-riding incentive of dominant emitters, either through reducing the costs of renewable energy or through enforceable mechanisms such as carbon border adjustments (Nordhaus, 2015). Once greed is addressed, fear-induced defection by vulnerable countries can be addressed through targeted climate finance and technology transfers (Okada, 2023). Ultimately, however, we require narratives that go beyond those of social dilemmas: models where sustainable behaviors are an equilibrium.
480 Our method offers a way to find them.

Code and data availability. Global CO₂ emissions data from fossil fuels and land-use change are from the Global Carbon Project (Friedlingstein et al., 2022), publicly available at <https://www.globalcarbonproject.org/>. RCP scenario data are from Moss et al. (2010). The model code, including the NumPyro/JAX implementation of all ten game configurations, the Bayesian inference pipeline, and the sensitivity analysis scripts, will be provided on request.



485 **Appendix A: Supplementary information overview**

The following sections provide detailed results for all model configurations. Appendix B and C present parameter estimates and fit diagnostics for the Cost-Benefit and Fear-Greed frameworks, respectively. Appendix D consolidates parameter sensitivity rankings across all models, including the two Chicken regimes. Appendix E reports RCP scenario alignment. Appendix F collects additional figures. Appendix G provides MCMC convergence diagnostics.

490 **Appendix B: Cost-Benefit framework**

Figures B1 and B2 show historical fit, strategy flow, sensitivity, and parameter posteriors for the Cost-Benefit models.

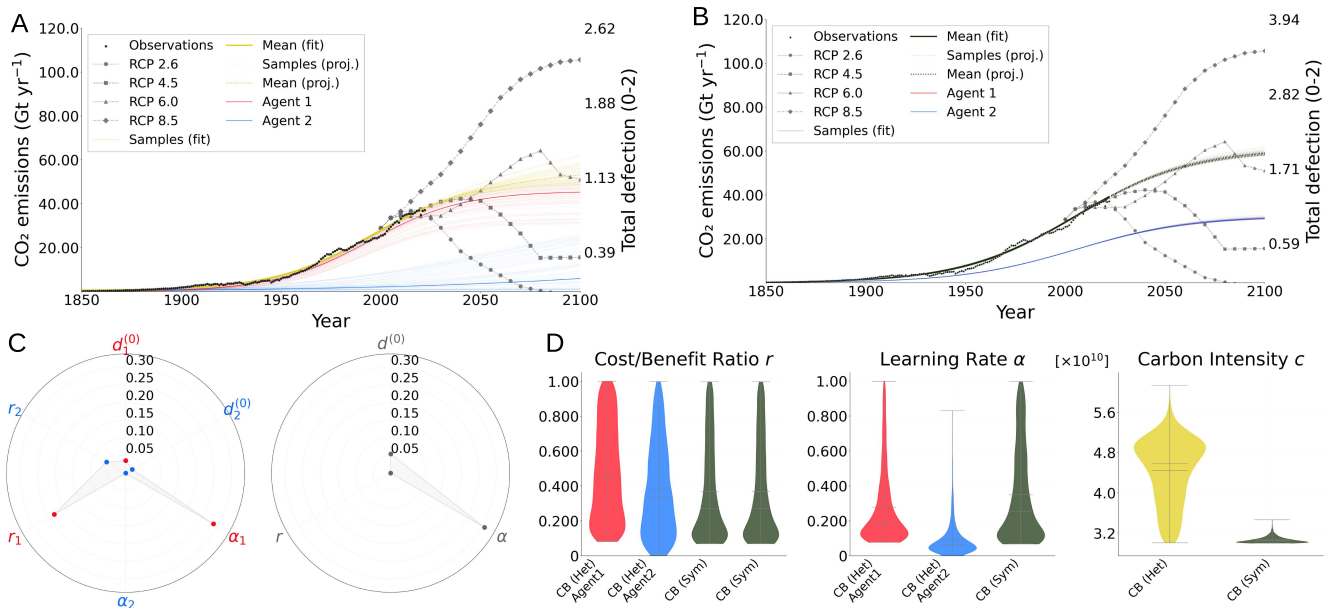


Figure B1. Cost-Benefit framework: heterogeneous vs. symmetric models. (a) Heterogeneous model (ELPD = 156.2, rank 2). Observed CO₂ emissions (black dots) with posterior fits and projections (yellow). Agent 1 (red) dominates cumulative emissions, while Agent 2 (blue) contributes marginally. Projections best align with RCP 6.0. (b) Symmetric model (ELPD = 150.7, rank 4, weight 0.05). Constraining agents to identical trajectories produces a smoother but underfitted historical fit. (c) Sensitivity analysis. Left: heterogeneous model; Agent 1’s learning rate (α_1 , 28.8%) and cost-benefit ratio (r_1 , 23.4%) are the primary leverage points, while α_2 has zero impact. Right: symmetric model; both shared adaptation speed (α , 30.7%) and cost-benefit ratio (r , 30.7%) equally drive emissions. (d) Posterior distributions. In the heterogeneous model, Agent 1 (red) shows higher r and faster learning than Agent 2 (blue). The symmetric model (dark green) recovers intermediate behavioral parameters but a lower carbon intensity (c), illustrating a trade-off between behavioral flexibility and physical scaling. Structurally a constrained Prisoner’s Dilemma ($f = g = r$), the heterogeneous CB variant produces predictions nearly redundant to the full PD ($\Delta\text{ELPD} = -1.9$).

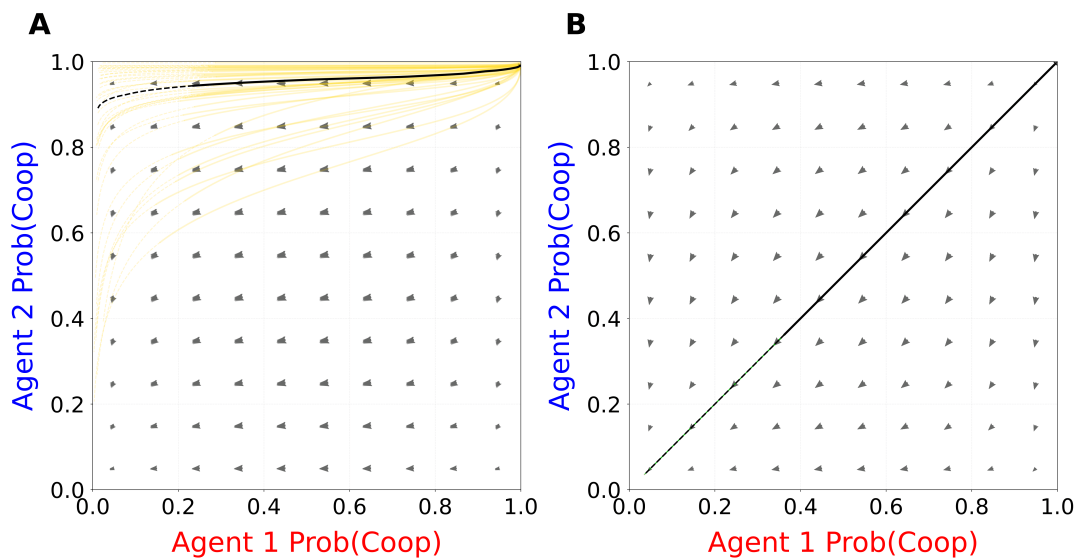


Figure B2. Phase space dynamics for the Cost-Benefit parameterization. Coupled strategy flows are shown for the Cost-Benefit model. The left panel illustrates a heterogeneous agent setup, while the right panel depicts a symmetric setup where trajectories are strictly confined to the main diagonal. Purple traces represent posterior simulations, with the solid black line indicating the median trajectory. Gray arrows depict the underlying gradient vector field.

Appendix C: Fear-Greed framework

Table C1 compares symmetric Fear-Greed games. PD and Chicken perform nearly identically (ELPD ≈ 151.7), while Stag Hunt and Harmony exhibit negative ELPD — structural failure. Posterior parameter estimates are given in Table C2.

Table C1. Symmetric Fear-Greed: ELPD and effective parameters.

Game	ELPD	p_{100}
Prisoner’s Dilemma	151.70	4.98
Chicken	151.68	5.27
Stag Hunt	−45.39	3.01
Harmony	−268.13	1.82



Table C2. Symmetric Fear-Greed: posterior means [95% HDI].

Parameter	PD	Chicken	Stag Hunt	Harmony
f	0.24 [0.00,0.63]	-0.55 [-1.00,-0.06]	0.62 [0.26,0.97]	-0.44 [-0.93,0.00]
g	0.54 [0.08,0.95]	0.31 [0.07,0.83]	-0.00 [-0.00,0.00]	-0.23 [-0.81,0.00]
α	0.21 [0.07,0.63]	0.39 [0.07,0.85]	0.46 [0.23,0.87]	0.01 [0.00,0.03]
d_0	0.003 [0.002,0.004]	0.001 [0.000,0.002]	0.055 [0.052,0.059]	0.065 [0.015,0.152]
c	3.65 [3.00,4.44]	11.62 [4.28,24.82]	3.03 [3.00,3.10]	12.19 [3.00,27.02]
σ	0.10 [0.09,0.11]	0.10 [0.09,0.11]	0.31 [0.28,0.35]	1.15 [1.03,1.27]

495 Table C3 quantifies ELPD gains from heterogeneity. PD gains +6.45, Stag Hunt gains +121 (but still inferior to PD). Chicken’s heterogeneous variant is bimodal and analyzed in Sect. 3.3. Posterior estimates for the heterogeneous PD (the best overall model) are shown in Table C4.

Table C3. ELPD improvement from heterogeneity.

Game	Symmetric	Heterogeneous	Δ ELPD
Prisoner’s Dilemma	151.70	158.15	+6.45
Chicken ^a	151.68	–	–
Stag Hunt	-45.39	75.62	+121.01
Harmony	-268.13	-268.48	-0.35

^a Bimodal posterior; regime-level ELPDs are 199.2 (R1) and 159.3 (R2). See Sect. 3.3.



Table C4. Heterogeneous PD: posterior estimates (mean, median, 50% and 95% HDI).

Parameter	Mean	Median	50% HDI	95% HDI
<i>Agent 1</i>				
f_1	0.384	0.334	[0.00,0.34]	[0.00,0.90]
g_1	0.395	0.339	[0.09,0.35]	[0.07,0.87]
α_1	0.333	0.270	[0.09,0.27]	[0.08,0.78]
$d_0^{(1)}$	0.002	0.002	[0.001,0.002]	[0.001,0.003]
<i>Agent 2</i>				
f_2	0.617	0.645	[0.64,1.00]	[0.17,1.00]
g_2	0.153	0.112	[0.01,0.12]	[0.00,0.42]
α_2	0.157	0.133	[0.09,0.16]	[0.00,0.36]
$d_0^{(2)}$	0.016	0.016	[0.01,0.02]	[0.01,0.03]
<i>Physical</i>				
c	3.605	3.353	[3.00,3.35]	[3.00,4.97]
σ	0.095	0.095	[0.09,0.10]	[0.09,0.11]

Appendix D: Parameter sensitivity analysis

This section consolidates sensitivity rankings across all model configurations. We perturbed each parameter by $\pm 20\%$ from its posterior mean and measured the resulting percentage change in cumulative 2023–2100 emissions (Eq. (9)). Values above 20% indicate the parameter amplifies emission outcomes; 20% indicates a proportional (linear) response. Agent 2’s learning rate shows moderate but secondary sensitivity (approximately 7–27% across models), substantially below Agent 1’s in PD and CB models but non-negligible.

D1 Sensitivity: Cost-Benefit framework

Table D1 reports sensitivity rankings for the Cost-Benefit framework. In the heterogeneous model, Agent 1’s learning rate dominates (28.8%), followed by the cost-to-benefit ratio r_1 (23.4%). The symmetric model shows that both the shared learning rate (α , 30.7%) and cost-benefit ratio (r , 30.7%) produce equally large responses, with carbon intensity (c) responding proportionally at 20.0%.



Table D1. Sensitivity analysis — Cost-Benefit framework. A 20% perturbation in each parameter produces the listed percentage change in cumulative 2023–2100 emissions (S_{cum}) and final-year emissions (S_{final}). Values above 20% indicate amplification.

Rank	Parameter	Input	Emission change (%)	
		Post. Mean	S_{cum}	S_{final}
<i>Heterogeneous</i>				
1	α_1	0.277	28.8%	10.6%
2	r_1	0.477	23.4%	3.3%
3	c	4.441	20.0%	20.0%
4	r_2	0.381	6.3%	7.3%
5	$d_0^{(1)}$	0.002	3.5%	0.1%
6	$d_0^{(2)}$	0.010	2.2%	1.9%
7	α_2	0.092	0.0%	0.0%
<i>Symmetric</i>				
1	α	0.349	30.7%	9.1%
1	r	0.344	30.7%	9.1%
3	c	3.063	20.0%	20.0%
4	d_0	0.004	5.4%	0.4%

D2 Sensitivity: Fear-Greed framework

510 Table D2 reports sensitivity rankings for all heterogeneous Fear-Greed models. Greed parameters consistently outweigh fear parameters: in the PD, a 20% perturbation in g_1 shifts emissions by 26.4%, whereas the same perturbation in f_1 produces only 1.0%. In Chicken, the pattern holds with g_1 at 23.1%–32.2% versus f_1 at 3.2%–7.9%. The heterogeneous cross-game radar comparison is presented in the main text (Fig. 4).



Table D2. Complete sensitivity rankings — heterogeneous Fear-Greed models. A 20% perturbation in each parameter produces the listed percentage change in cumulative (S_{cum}) and final-year (S_{final}) emissions. Values above 20% indicate the parameter amplifies emission outcomes. For α_2 , the perturbation is applied via the learning rate ratio $\rho = \alpha_2/\alpha_1$ (equivalent to a 20% change in α_2 at fixed α_1).

Rank	Parameter	Input Post. Mean	Emission change (%)	
			S_{cum}	S_{final}
<i>PD (Het.)</i>				
1	α_1	0.333	32.0%	12.8%
2	g_1	0.395	26.4%	6.4%
3	c	3.605	20.0%	20.1%
4	α_2	0.157	7.9%	5.1%
5	f_2	0.617	5.4%	3.7%
6	$d_0^{(1)}$	0.002	3.8%	0.4%
7	$d_0^{(2)}$	0.016	2.4%	0.7%
8	g_2	0.153	2.3%	0.9%
9	f_1	0.384	1.0%	0.2%
<i>Stag Hunt (Het.)</i>				
1	α_1	0.991	29.3%	13.0%
2	$d_0^{(2)}$	0.070	24.9%	6.0%
3	f_1	0.991	21.9%	4.6%
4	c	3.021	20.0%	20.1%
5	α_2	0.311	8.4%	5.3%
5	f_2	0.404	8.4%	5.3%
7	$d_0^{(1)}$	0.006	5.2%	0.7%
8	g_1	-0.001	0.2%	0.1%
9	g_2	-0.000	0.1%	0.1%

D3 Chicken game: regime sensitivity

515 The bimodality of the heterogeneous Chicken game, its two regimes, and their structural resemblance to the Prisoner’s Dilemma are analyzed in Sect. 3.3 (Table 3, Fig. 3). Table D3 below reports the sensitivity rankings for the two regimes. In both, greed parameters and Agent 1’s greed dominate, while fear parameters have negligible influence. The α_1 sensitivity collapses to near zero in both regimes because Agent 1 reaches a corner solution ($d_1 \approx 1$) early in the projection horizon.



Table D3. Sensitivity rankings for heterogeneous Chicken regimes. A 20% perturbation in each parameter produces the listed percentage change in cumulative emissions. Agent 1’s learning rate has zero sensitivity in both regimes. Only the top six parameters are shown.

Regime 1 (fast A1, clean economy)			Regime 2 (moderate A1, dirty economy)		
Rank	Parameter	S_{cum}	Rank	Parameter	S_{cum}
1	g_1	23.1%	1	g_1	32.2%
2	α_2	20.6%	2	α_2	26.9%
3	c	20.0%	3	c	20.0%
4	$d_0^{(2)}$	15.2%	4	$d_0^{(2)}$	9.1%
5	g_2	10.3%	5	g_2	7.2%
6	f_1	7.9%	6	$d_0^{(1)}$	4.4%

Appendix E: RCP scenario alignment

520 Table E1 reports the three best-matching RCP scenarios for each model. All social dilemma models align most closely with RCP 6.0 (RMSE 5.6–11.4 Gt CO₂/yr), except Chicken Regime 1 which aligns with RCP 4.5. Figure E1 displays the full alignment matrix.

Table E1. Top three RCP matches per model (RMSE, Gt CO₂/yr). The two heterogeneous Chicken regimes are reported separately (Sect. 3.3); the aggregate heterogeneous Chicken row is omitted because its bimodal posterior conflates the two regimes.

Model	1st	RMSE	2nd	RMSE	3rd	RMSE
Cost-Benefit (Sym.)	RCP 6.0	5.56	RCP 4.5	26.99	RCP 8.5	34.38
Stag Hunt (Het.)	RCP 6.0	5.58	RCP 4.5	24.65	RCP 8.5	37.26
Chicken (Sym.)	RCP 6.0	5.69	RCP 4.5	26.88	RCP 8.5	34.67
PD (Sym.)	RCP 6.0	5.78	RCP 4.5	28.05	RCP 8.5	33.51
PD (Het.)	RCP 6.0	6.37	RCP 4.5	28.13	RCP 8.5	33.49
Cost-Benefit (Het.)	RCP 6.0	6.95	RCP 4.5	22.37	RCP 8.5	39.35
Chicken (Regime 2)	RCP 6.0	9.28	RCP 4.5	20.80	RCP 2.6	38.52
Stag Hunt (Sym.)	RCP 6.0	11.38	RCP 8.5	29.81	RCP 4.5	31.25
Chicken (Regime 1)	RCP 4.5	12.36	RCP 6.0	16.99	RCP 2.6	27.89

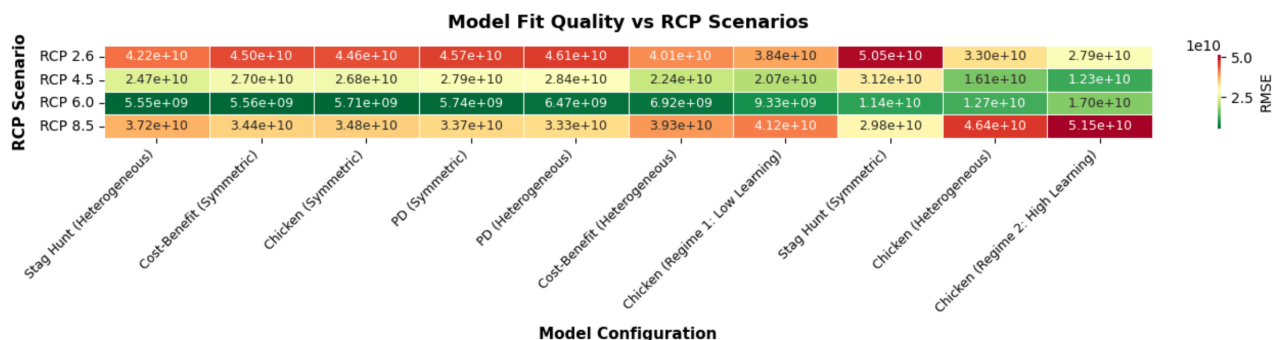


Figure E1. Model–RCP alignment heatmap. Cool colors indicate lower RMSE (better alignment). Most social dilemma models cluster around RCP 6.0; the historically implausible Chicken Regime 1 is the only exception, aligning with RCP 4.5.



Appendix F: Additional figures

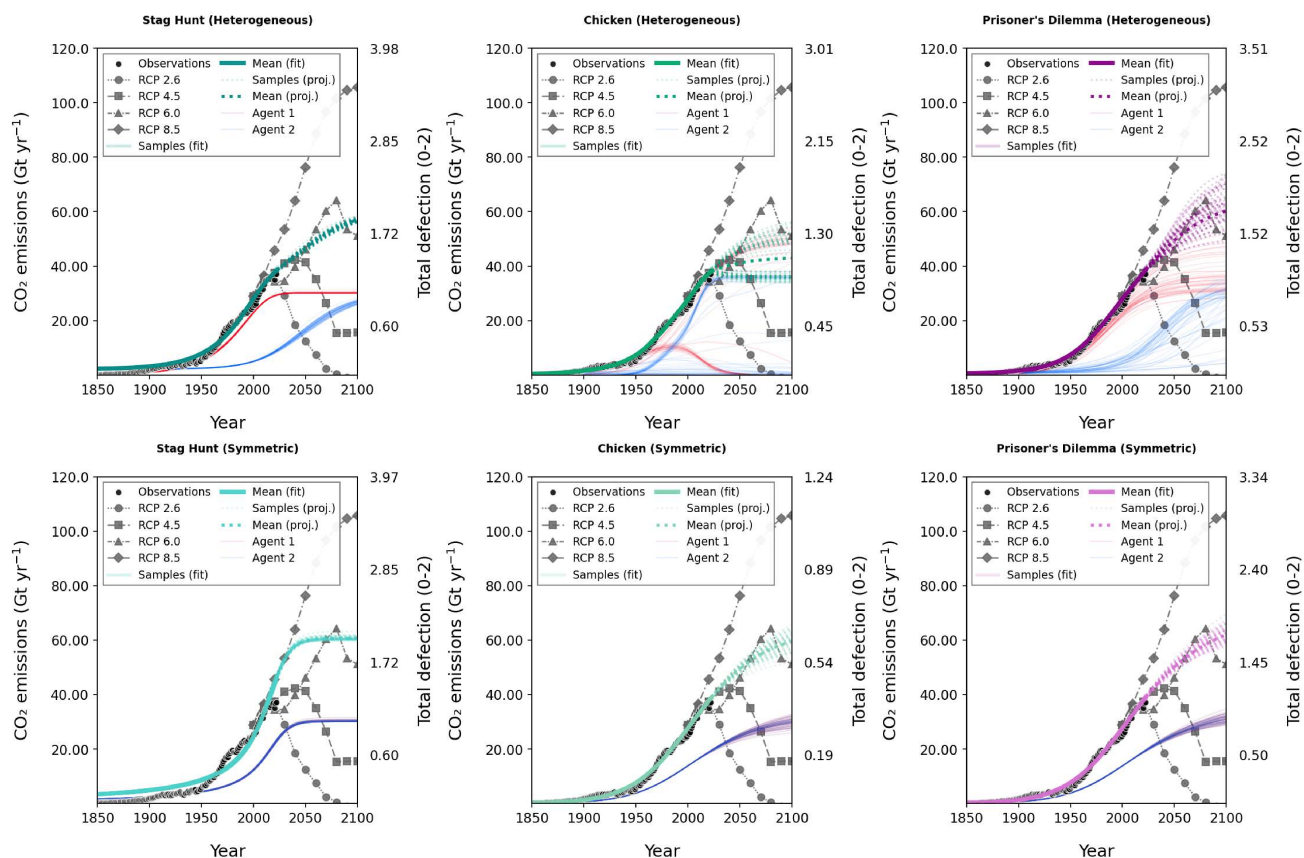


Figure F1. Fear-Greed model projections. Historical fits and 95% HDI for heterogeneous (top) and symmetric (bottom) configurations across Stag Hunt (SH), Chicken (CH), and Prisoner's Dilemma (PD). In the heterogeneous models (top), Agent 1 (red) and Agent 2 (blue) trajectories are distinct. In the symmetric models (bottom), the mean fits are color-coded by game type (SH: yellow, CH: cyan, PD: orange) alongside the Agent 2 trajectory (blue).

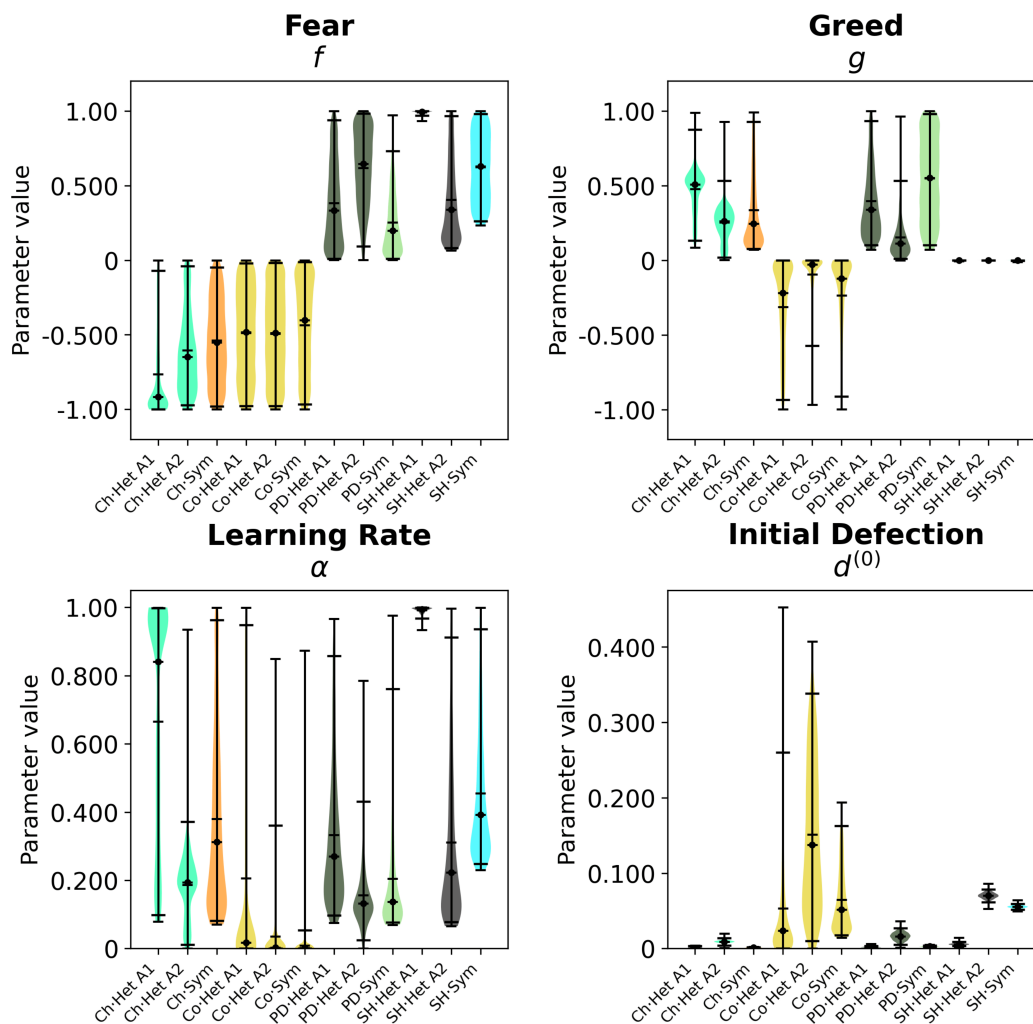


Figure F2. Fear-Greed parameter posteriors. Distributions of f , g , α , and d_0 across game types and agent configurations.

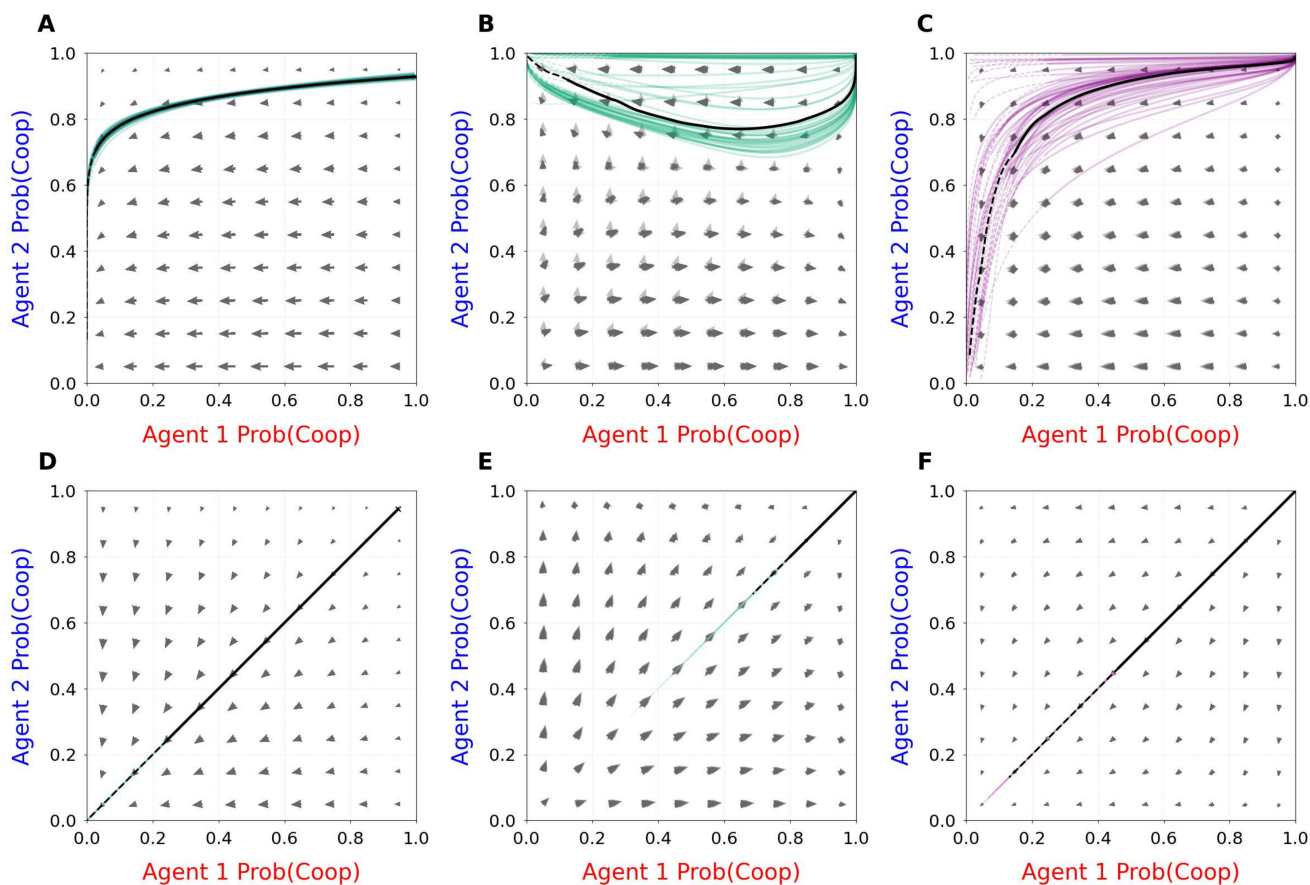


Figure F3. Phase space dynamics in the Fear-Greed parameter space. Coupled strategy flows are shown for the Stag Hunt (a, d), Chicken (b, e), and Prisoner's Dilemma (c, f). The top row (a–c) illustrates heterogeneous agent setups, while the bottom row (d–f) depicts symmetric setups where trajectories converge along the main diagonal. Colored traces (teal for Stag Hunt, green for Chicken, and purple for Prisoner's Dilemma) represent posterior simulations for each respective game, with the solid black line indicating the median trajectory. Gray arrows depict the underlying gradient vector field.

Appendix G: MCMC diagnostics

525 Tables G1–G3 report convergence diagnostics for all model configurations. Note the severe non-convergence ($\hat{R} \gg 1.1$) in the raw heterogeneous Chicken model, which motivated the regime-level analysis in Sect. 3.3.



Table G1. MCMC diagnostics for the Cost-Benefit framework.

Parameter	ESS	\hat{R}
<i>Symmetric</i>		
Emission scale (c)	425	1.01
Learning rate (α)	171	1.02
Cost/benefit (r)	171	1.02
Initial defection (d_0)	613	1.00
Noise level (σ)	398	1.01
<i>Heterogeneous</i>		
Emission scale (c)	173	1.01
Learning rate Agt 1 (α_1)	249	1.02
Learning rate Agt 2 (α_2)	357	1.01
Cost/benefit Agt 1 (r_1)	241	1.02
Cost/benefit Agt 2 (r_2)	425	1.01
Initial defection Agt 1 ($d_0^{(1)}$)	166	1.02
Initial defection Agt 2 ($d_0^{(2)}$)	257	1.02
Noise level (σ)	600	1.01



Table G2. MCMC diagnostics for all Fear-Greed games. Values for heterogeneous models are reported as *Agent 1 / Agent 2*. Note the severe non-convergence ($\hat{R} \gg 1.1$) in the raw heterogeneous Chicken model.

	Symmetric		Heterogeneous	
	ESS	\hat{R}	ESS	\hat{R}
Prisoner's Dilemma				
Learning rate (α)	251	1.01	457 / 420	1.01 / 1.01
Fear (f)	299	1.00	774 / 686	1.00 / 1.00
Greed (g)	248	1.01	451 / 519	1.01 / 1.00
Initial defection (d_0)	461	1.00	598 / 566	1.01 / 1.01
Emission scale (c)	392	1.00	385	1.01
Noise level (σ)	673	1.00	1062	1.00
Stag Hunt				
Learning rate (α)	1614	1.00	3071 / 1258	1.00 / 1.00
Fear (f)	1624	1.00	2951 / 1215	1.00 / 1.00
Greed (g)	1719	1.00	3285 / 1856	1.00 / 1.00
Initial defection (d_0)	2246	1.00	1860 / 1807	1.00 / 1.00
Emission scale (c)	1819	1.00	3187	1.00
Noise level (σ)	2174	1.00	3674	1.00
Harmony				
Learning rate (α)	1372	1.00	413 / 792	1.02 / 1.01
Fear (f)	2552	1.00	3996 / 3424	1.00 / 1.00
Greed (g)	1681	1.00	1080 / 801	1.00 / 1.01
Initial defection (d_0)	1228	1.00	1373 / 886	1.00 / 1.01
Emission scale (c)	1199	1.00	2689	1.00
Noise level (σ)	2447	1.00	3572	1.00
Chicken				
Learning rate (α)	207	1.00	3 / 3	1.81 / 1.83
Fear (f)	290	1.01	7 / 3	1.55 / 1.67
Greed (g)	204	1.00	3 / 57	1.83 / 1.43
Initial defection (d_0)	156	1.01	3 / 3	1.83 / 1.83
Emission scale (c)	153	1.01	3	1.83
Noise level (σ)	579	1.00	3	1.82



Table G3. Detailed diagnostics for the heterogeneous Chicken game. Clustering the raw posterior reveals two distinct, well-converged regimes (see Sect. 3.3).

	Regime 1		Regime 2	
	<i>(Fast AI, clean economy)</i>		<i>(Moderate AI, dirty economy)</i>	
	ESS	\hat{R}	ESS	\hat{R}
Learning rate Agent 1 (α_1)	492	1.011	148	1.024
Learning rate Agent 2 (α_2)	154	1.019	186	1.013
Fear Agent 1 (f_1)	855	1.002	271	1.004
Fear Agent 2 (f_2)	165	1.017	717	1.007
Greed Agent 1 (g_1)	539	1.011	508	1.006
Greed Agent 2 (g_2)	170	1.014	656	1.010
Initial defection Agent 1 ($d_0^{(1)}$)	507	1.009	140	1.012
Initial defection Agent 2 ($d_0^{(2)}$)	679	1.008	794	1.002
Emission scale (c)	769	1.008	610	1.003
Noise level (σ)	1157	1.004	504	1.004

Author contributions. A.B. and W.B. conceptualized the study. A.B. built the game-theoretic model, implemented the Bayesian probabilistic programming methodology, conducted all analyses, and wrote the first draft of the manuscript. H.S. contributed to the Bayesian methodology and provided supervision. W.B. supervised the project throughout. All authors wrote, reviewed, and edited the manuscript.

530 *Competing interests.* The authors declare that they have no conflict of interest.



References

- Ahn, T.-K., Ostrom, E., Schmidt, D., Shupp, R., and Walker, J.: Cooperation in PD games: Fear, greed, and history of play, *Public Choice*, 106, 137–155, <https://doi.org/10.1023/A:1005219123532>, 2001.
- Aklin, M. and Mildenerger, M.: Prisoners of the wrong dilemma: Why distributive conflict, not collective action, characterizes the politics
535 of climate change, *Global Environmental Politics*, 20, 4–27, https://doi.org/10.1162/glep_a_00578, 2020.
- Armstrong McKay, D. I., Staal, A., Abrams, J. F., Winkelmann, R., Sakschewski, B., Loriani, S., Fetzer, I., Cornell, S. E., Rockström, J., and Lenton, T. M.: Exceeding 1.5°C global warming could trigger multiple climate tipping points, *Science*, 377, eabn7950, <https://doi.org/10.1126/science.abn7950>, 2022.
- Barde, S.: Direct comparison of agent-based models of herding in financial markets, *Journal of Economic Dynamics and Control*, 73, 329–
540 353, <https://doi.org/10.1016/j.jedc.2016.10.005>, 2016.
- Barfuss, W. and Mann, R. P.: Modeling the effects of environmental and perceptual uncertainty using deterministic reinforcement learning dynamics with partial observability, *Physical Review E*, 105, 034 409, <https://doi.org/10.1103/PhysRevE.105.034409>, 2022.
- Barfuss, W., Donges, J. F., Lade, S. J., and Kurths, J.: When optimization for governing human-environment tipping elements is neither sustainable nor safe, *Nature Communications*, 9, 2354, <https://doi.org/10.1038/s41467-018-04738-z>, 2018.
- 545 Barfuss, W., Donges, J. F., and Kurths, J.: Deterministic limit of temporal difference reinforcement learning for stochastic games, *Physical Review E*, 99, 043 305, <https://doi.org/10.1103/PhysRevE.99.043305>, 2019.
- Barfuss, W., Donges, J. F., Wiedermann, M., and Lucht, W.: Caring for the future can turn tragedy into comedy for long-term collective action under risk of collapse, *Proceedings of the National Academy of Sciences*, 117, 12 915–12 922, <https://doi.org/10.1073/pnas.1916545117>, 2020.
- 550 Barrett, S.: *Environment and Statecraft: The Strategy of Environmental Treaty-Making*, Oxford University Press, Oxford, 2003.
- Barrett, S.: *Why Cooperate? The Incentive to Supply Global Public Goods*, Oxford University Press, Oxford, 2007.
- Bijak, J.: *Towards Bayesian Model-Based Demography: Agency, Complexity and Uncertainty in Migration Studies*, Methodos Series, Springer, <https://doi.org/10.1007/978-3-030-83039-7>, 2021.
- Bloembergen, D., Tuyls, K., Hennes, D., and Kaisers, M.: Evolutionary dynamics of multi-agent learning: A survey, *Journal of Artificial
555 Intelligence Research*, 53, 659–697, <https://doi.org/10.1613/jair.4818>, 2015.
- Böhringer, C., Carbone, J. C., and Rutherford, T. F.: The Strategic Value of Carbon Tariffs, *American Economic Journal: Economic Policy*, 8, 28–51, <https://doi.org/10.1257/pol.20130327>, 2016.
- Börgers, T. and Sarin, R.: Learning through reinforcement and replicator dynamics, *Journal of Economic Theory*, 77, 1–14, <https://doi.org/10.1006/jeth.1997.2319>, 1997.
- 560 Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q.: JAX: Composable transformations of Python+NumPy programs, <http://github.com/jax-ml/jax>, version 0.3.13, 2018.
- Buchholz, W. and Sandler, T.: Global public goods: A survey, *Journal of Economic Literature*, 59, 488–545, <https://doi.org/10.1257/jel.20190696>, 2021.
- Camerer, C. F.: *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press, Princeton, NJ, 2003.
- 565 Carrozzo Magli, A. and Manfredi, P.: Coordination games vs prisoner’s dilemma in sustainability games: A critique of recent contributions and a discussion of policy implications, *Ecological Economics*, 192, 107 264, <https://doi.org/10.1016/j.ecolecon.2021.107264>, 2022.



- Coombs, C. H.: A reparameterization of the prisoner's dilemma game, *Behavioral Science*, 18, 424–428, <https://doi.org/10.1002/bs.3830180605>, 1973.
- Crawford, V. P., Costa-Gomes, M. A., and Iriberry, N.: Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications, *Journal of Economic Literature*, 51, 5–62, <https://doi.org/10.1257/jel.51.1.5>, 2013.
- 570 Crutzen, P. J.: Geology of mankind, *Nature*, 415, 23, <https://doi.org/10.1038/415023a>, 2002.
- DeCanio, S. J. and Fremstad, A.: Game Theory and Climate Diplomacy, *Ecological Economics*, 85, 177–187, <https://doi.org/10.1016/j.ecolecon.2011.04.016>, 2013.
- Donges, J. F. and Barfuss, W.: From math to metaphors and back again: Social-ecological resilience from a multi-agent-environment perspective, *GAIA – Ecological Perspectives for Science and Society*, 26, 182–190, <https://doi.org/10.14512/gaia.26.S1.5>, 2017.
- 575 Donges, J. F., Winkelmann, R., Lucht, W., Cornell, S. E., Dyke, J. G., Rockström, J., Heitzig, J., and Schellnhuber, H. J.: Closing the loop: Reconnecting human dynamics to Earth System science, *The Anthropocene Review*, 4, 151–157, <https://doi.org/10.1177/2053019617725537>, 2017.
- Donges, J. F., Heitzig, J., Barfuss, W., Wiedermann, M., Kassel, J. A., Kittel, T., Kolb, J. J., Kolster, T., Müller-Hansen, F., Otto, I. M., Zimmerer, K. B., and Lucht, W.: Earth system modeling with endogenous and dynamic human societies: the copan: CORE open World–Earth modeling framework, *Earth System Dynamics*, 11, 395–413, <https://doi.org/10.5194/esd-11-395-2020>, 2020.
- 580 Dorninger, C., Hornborg, A., Abson, D. J., von Wehrden, H., Schaffartzik, A., Giljum, S., Engler, J.-O., Feller, R. L., Hubacek, K., and Wieland, H.: Global patterns of ecologically unequal exchange: Implications for sustainability in the 21st century, *Ecological Economics*, 179, 106 824, <https://doi.org/10.1016/j.ecolecon.2020.106824>, 2021.
- 585 Dutton, J. M. and Starbuck, W. H.: *Computer Simulation of Human Behavior*, Wiley, New York, 1971.
- Frederick, S., Loewenstein, G., and O'Donoghue, T.: Time Discounting and Time Preference: A Critical Review, *Journal of Economic Literature*, 40, 351–401, <https://doi.org/10.1257/002205102320161311>, 2002.
- Friedlingstein, P., O'Sullivan, M., Jones, M. W., Andrew, R. M., Gregor, L., Hauck, J., Le Quéré, C., Luijkx, I. T., Olsen, A., Peters, G. P., et al.: Global Carbon Budget 2022, *Earth System Science Data*, 14, 4811–4900, <https://doi.org/10.5194/essd-14-4811-2022>, 2022.
- 590 Friedman, M.: *The Methodology of Positive Economics*, in: *Essays in Positive Economics*, pp. 3–43, University of Chicago Press, Chicago, 1953.
- Fudenberg, D. and Levine, D. K.: Whither game theory? Towards a theory of learning in games, *Journal of Economic Perspectives*, 30, 151–170, <https://doi.org/10.1257/jep.30.4.151>, 2016.
- García, J. and Traulsen, A.: Picking strategies in games of cooperation, *Proceedings of the National Academy of Sciences*, 122, <https://doi.org/10.1073/pnas.2319925121>, 2025.
- 595 Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B.: *Bayesian Data Analysis*, Chapman and Hall/CRC, 3rd edn., <https://doi.org/10.1201/b16018>, 2013.
- Gifford, R.: The dragons of inaction: Psychological barriers that limit climate change mitigation and adaptation, *American Psychologist*, 66, 290–302, <https://doi.org/10.1037/a0023566>, 2011.
- 600 Grazzini, J., Richiardi, M. G., and Tsionas, M.: Bayesian estimation of agent-based models, *Journal of Economic Dynamics and Control*, 77, 26–47, <https://doi.org/10.1016/j.jedc.2017.01.014>, 2017.
- Guerini, M. and Moneta, A.: A method for agent-based models validation, *Journal of Economic Dynamics and Control*, 82, 125–141, <https://doi.org/10.1016/j.jedc.2017.06.001>, 2017.
- Hale, T.: Catalytic cooperation, *Global Environmental Politics*, 20, 73–98, https://doi.org/10.1162/glep_a_00561, 2020.



- 605 Hardin, G.: The Tragedy of the Commons, *Science*, 162, 1243–1248, <https://doi.org/10.1126/science.162.3859.1243>, 1968.
- Hartig, F., Calabrese, J. M., Reineking, B., Wiegand, T., and Huth, A.: Statistical inference for stochastic simulation models – theory and application, *Ecology Letters*, 14, 816–827, <https://doi.org/10.1111/j.1461-0248.2011.01640.x>, 2011.
- Hickel, J.: Quantifying national responsibility for climate breakdown: an equality-based attribution approach for carbon dioxide emissions in excess of the planetary boundary, *The Lancet Planetary Health*, 4, e399–e404, [https://doi.org/10.1016/S2542-5196\(20\)30196-0](https://doi.org/10.1016/S2542-5196(20)30196-0), 2020.
- 610 Hilbe, C., Nowak, M. A., and Sigmund, K.: Evolution of extortion in iterated prisoner’s dilemma games, *Proceedings of the National Academy of Sciences*, 110, 6913–6918, <https://doi.org/10.1073/pnas.1214834110>, 2013.
- Hofbauer, J. and Sigmund, K.: *Evolutionary Games and Population Dynamics*, Cambridge University Press, Cambridge, 1998.
- Keppo, I., Butnar, I., Bauer, N., Caspani, M., Edelenbosch, O., Emmerling, J., Fragkos, P., Guivarch, C., Harmsen, M., Lefèvre, J., Le Gallouët, T., Markandya, A., Mima, S., Perdana, S., Peters, M., Petropoulos, A., Reanos, M. A. T., Schaeffer, R., Shiogama, H., Vrontisi, Z.,
- 615 Fragiadakis, K., Weyant, J., and van Vuuren, D. P.: Exploring the possibility space: taking stock of the diverse capabilities and gaps in integrated assessment models, *Environmental Research Letters*, 16, 053 006, <https://doi.org/10.1088/1748-9326/abe5d8>, 2021.
- Klinsky, S., Roberts, T., Huber, S., Knutti, R., King, A. D., Korze, D. L., Mace, M., Mahon, R., Shue, H., Steinberger, J. K., and Sagar, A. D.: Why equity is fundamental in climate change policy research, *Global Environmental Change*, 44, 170–173, <https://doi.org/10.1016/j.gloenvcha.2016.08.002>, 2017.
- 620 Kollmuss, A. and Agyeman, J.: Mind the Gap: Why do people act environmentally and what are the barriers to pro-environmental behavior?, *Environmental Education Research*, 8, 239–260, <https://doi.org/10.1080/13504620220145401>, 2002.
- Kollock, P.: Social dilemmas: The anatomy of cooperation, *Annual Review of Sociology*, 24, 183–214, <https://doi.org/10.1146/annurev.soc.24.1.183>, 1998.
- Krapu, C. and Borsuk, M.: Probabilistic programming: A review for environmental modellers, *Environmental Modelling & Software*, 114, 40–48, <https://doi.org/10.1016/j.envsoft.2019.01.014>, 2019.
- 625 Kuleshov, V. and Schrijvers, O.: Inverse game theory: Learning utilities in succinct games, in: *Proceedings of the 11th International Conference on Web and Internet Economics (WINE)*, pp. 413–427, Springer, Berlin, Heidelberg, https://doi.org/10.1007/978-3-662-48995-6_30, 2015.
- Kumar, R., Carroll, C., Hartikainen, A., and Martin, O.: ArviZ a unified library for exploratory analysis of Bayesian models in Python, *Journal of Open Source Software*, 4, 1143, <https://doi.org/10.21105/joss.01143>, 2019.
- 630 Lamperti, F., Roventini, A., and Sani, A.: Agent-based model calibration using machine learning surrogates, *Journal of Economic Dynamics and Control*, 90, 366–389, <https://doi.org/10.1016/j.jedc.2018.03.011>, 2018.
- Lange, A., Vogt, C., and Ziegler, A.: On the importance of equity in international climate policy: An empirical analysis, *Energy Economics*, 29, 545–562, <https://doi.org/10.1016/j.eneco.2006.09.002>, 2007.
- 635 Lenton, T. M., Held, H., Kriegler, E., Hall, J. W., Lucht, W., Rahmstorf, S., and Schellnhuber, H. J.: Tipping elements in the Earth’s climate system, *Proceedings of the National Academy of Sciences*, 105, 1786–1793, <https://doi.org/10.1073/pnas.0705414105>, 2008.
- Lessmann, K., Marschinski, R., and Edenhofer, O.: The effects of tariffs on coalition formation in a dynamic global warming game, *Economic Modelling*, 26, 641–649, <https://doi.org/10.1016/j.econmod.2009.01.005>, 2009.
- Levin, K., Cashore, B., Bernstein, S., and Auld, G.: Overcoming the tragedy of super wicked problems: constraining our future selves to
- 640 ameliorate global climate change, *Policy Sciences*, 45, 123–152, <https://doi.org/10.1007/s11077-012-9151-0>, 2012.
- Levins, R.: The strategy of model building in population biology, *American Scientist*, 54, 421–431, 1966.



- Lipman, B. L.: Cooperation among egoists in Prisoners' Dilemma and Chicken games, *Public Choice*, 51, 315–331, <https://doi.org/10.1007/BF00128881>, 1986.
- Macy, M. W. and Flache, A.: Learning dynamics in social dilemmas, *Proceedings of the National Academy of Sciences*, 99, 7229–7236, <https://doi.org/10.1073/pnas.092080099>, 2002.
- 645 Markowitz, E. M. and Shariff, A. F.: Climate change and moral judgement, *Nature Climate Change*, 2, 243–247, <https://doi.org/10.1038/nclimate1378>, 2012.
- McKelvey, R. D. and Palfrey, T. R.: Quantal response equilibria for normal form games, *Games and Economic Behavior*, 10, 6–38, <https://doi.org/10.1006/game.1995.1023>, 1995.
- 650 Moser, S. C.: Communicating climate change: history, challenges, process and future directions, *WIREs Climate Change*, 1, 31–53, <https://doi.org/10.1002/wcc.11>, 2010.
- Moss, R. H., Edmonds, J. A., Hibbard, K. A., Manning, M. R., Rose, S. K., van Vuuren, D. P., Carter, T. R., Emori, S., Kainuma, M., Kram, T., Meehl, G. A., Mitchell, J. F. B., Nakicenovic, N., Riahi, K., Smith, S. J., Stouffer, R. J., Thomson, A. M., Weyant, J. P., and Wilbanks, T. J.: The next generation of scenarios for climate change research and assessment, *Nature*, 463, 747–756, <https://doi.org/10.1038/nature08823>, 2010.
- 655 Nordhaus, W.: Climate clubs: Overcoming free-riding in international climate policy, *American Economic Review*, 105, 1339–1370, <https://doi.org/10.1257/aer.15000001>, 2015.
- Okada, A.: Cooperation through transfers in a prisoner's dilemma, *Games and Economic Behavior*, 141, 428–448, <https://doi.org/10.1016/j.geb.2023.07.006>, 2023.
- 660 Okereke, C.: Equity norms in global environmental governance, *Global Environmental Politics*, 8, 25–50, <https://doi.org/10.1162/glep.2008.8.3.25>, 2008.
- Ostrom, E.: *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge University Press, Cambridge, 1990.
- Otto, I. M., Donges, J. F., Cremades, R., Bhowmik, A., Hewitt, R. J., Lucht, W., Rockström, J., Allerberger, F., McCaffrey, M., Doe, S. P. A., Lenferna, A., Morán, N., van Vuuren, D. P., and Schellnhuber, H. J.: Social tipping dynamics for stabilizing Earth's climate by 2050, *Proceedings of the National Academy of Sciences*, 117, 2354–2365, <https://doi.org/10.1073/pnas.1900577117>, 2020.
- 665 Ouellette, J. A. and Wood, W.: Habit and intention in everyday life: The multiple processes by which past behavior predicts future behavior, *Psychological Bulletin*, 124, 54–74, <https://doi.org/10.1037/0033-2909.124.1.54>, 1998.
- Paavola, J. and Adger, W. N.: Fair adaptation to climate change, *Ecological Economics*, 56, 594–609, <https://doi.org/10.1016/j.ecolecon.2005.03.015>, 2006.
- 670 Phan, D., Pradhan, N., and Jankowiak, M.: Composable effects for flexible and accelerated probabilistic programming in NumPyro, arXiv preprint arXiv:1912.11554, 2019.
- Radosavljevic, S., Haider, L. J., Lade, S. J., and Schlüter, M.: Dynamical systems modeling for structural understanding of social-ecological systems: A primer, *Ecological Complexity*, 56, 101 052, <https://doi.org/10.1016/j.ecocom.2023.101052>, 2023.
- Richardson, K., Steffen, W., Lucht, W., Bendtsen, J., Cornell, S. E., Donges, J. F., Drüke, M., Fetzer, I., Bala, G., von Bloh, W., Feulner, G., Fiedler, S., Gerten, D., Gleeson, T., Hofmann, M., Huiskamp, W., Kummu, M., Moberg, C., Nogues-Bravo, D., Pedde, S., Popp, A., Reyer, C., Rockström, J., Scheffer, M., Schellnhuber, H. J., Schlüter, M., Sgubin, G., Snyder, A., Springmann, M., and Walker, B.: Earth beyond six of nine planetary boundaries, *Science Advances*, 9, eadh2458, <https://doi.org/10.1126/sciadv.adh2458>, 2023.
- Rockström, J., Steffen, W., Noone, K., Persson, Å., Chapin, F. S., Lambin, E. F., Lenton, T. M., Scheffer, M., Folke, C., Schellnhuber, H. J., Nykvist, B., de Wit, C. A., Hughes, T., van der Leeuw, S., Rodhe, H., Sörlin, S., Snyder, P. K., Costanza, R., Svedin, U., Falkenmark,



- 680 M., Karlberg, L., Corell, R. W., Fabry, V. J., Hansen, J., Walker, B., Liverman, D., Richardson, K., Crutzen, P., and Foley, J. A.: A safe operating space for humanity, *Nature*, 461, 472–475, <https://doi.org/10.1038/461472a>, 2009.
- Santos, F. C., Vasconcelos, V. V., Santos, M. D., Neves, P. N. B., and Pacheco, J. M.: Evolutionary dynamics of climate change under collective-risk dilemmas, *Mathematical Models and Methods in Applied Sciences*, 22, 1140004, <https://doi.org/10.1142/S0218202511400045>, 2012.
- 685 Schinko, T., Mechler, R., and Hochrainer-Stigler, S.: A methodological framework to operationalize climate risk management: managing sovereign climate-related extreme event risk in Austria, in: *Loss and Damage from Climate Change: Concepts, Methods and Policy Options*, edited by Mechler, R., Bouwer, L. M., Schinko, T., Surminski, S., and Linnerooth-Bayer, J., pp. 299–322, Springer, https://doi.org/10.1007/978-3-319-72026-5_12, 2019.
- Schulze, J., Müller, B., Groeneveld, J., and Grimm, V.: Agent-based modelling of social-ecological systems: Achievements, challenges, and
690 a way forward, *Journal of Artificial Societies and Social Simulation*, 20, 8, <https://doi.org/10.18564/jasss.3423>, 2017.
- Simon, H. A.: A behavioral model of rational choice, *The Quarterly Journal of Economics*, 69, 99–118, <https://doi.org/10.2307/1884852>, 1955.
- Skyrms, B.: *The Stag Hunt and the Evolution of Social Structure*, Cambridge University Press, Cambridge, 2004.
- Solomon, S., Plattner, G.-K., Knutti, R., and Friedlingstein, P.: Irreversible climate change due to carbon dioxide emissions, *Proceedings of
695 the National Academy of Sciences*, 106, 1704–1709, <https://doi.org/10.1073/pnas.0812721106>, 2009.
- Steffen, W., Broadgate, W., Deutsch, L., Gaffney, O., and Ludwig, C.: The trajectory of the Anthropocene: The Great Acceleration, *The Anthropocene Review*, 2, 81–98, <https://doi.org/10.1177/2053019614564785>, 2015a.
- Steffen, W., Richardson, K., Rockström, J., Cornell, S. E., Fetzer, I., Bennett, E. M., Biggs, R., Carpenter, S. R., de Vries, W., de Wit, C. A.,
Folke, C., Gerten, D., Heinke, J., Mace, G. M., Persson, L. M., Ramanathan, V., Reyers, B., and Sörlin, S.: Planetary boundaries: Guiding
700 human development on a changing planet, *Science*, 347, 1259–1265, <https://doi.org/10.1126/science.1259855>, 2015b.
- Steffen, W., Rockström, J., Richardson, K., Lenton, T. M., Folke, C., Liverman, D., Summerhayes, C. P., Barnosky, A. D., Cornell, S. E.,
Crucifix, M., Donges, J. F., Fetzer, I., Lade, S. J., Scheffer, M., Winkelmann, R., and Schellnhuber, H. J.: Trajectories of the Earth System
in the Anthropocene, *Proceedings of the National Academy of Sciences*, 115, 8252–8259, <https://doi.org/10.1073/pnas.1810141115>, 2018.
- Stoknes, P. E.: Rethinking climate communications and the “psychological climate paradox”, *Energy Research & Social Science*, 1, 161–170,
705 <https://doi.org/10.1016/j.erss.2014.03.007>, 2014.
- Storm, H., Heckelei, T., and Baylis, K.: Probabilistic programming for embedding theory and quantifying uncertainty in econometric analysis,
European Review of Agricultural Economics, 51, 589–616, <https://doi.org/10.1093/erae/jbae016>, 2024.
- Tilman, A. R., Plotkin, J. B., and Akçay, E.: Evolutionary games with environmental feedbacks, *Nature Communications*, 11, 915,
<https://doi.org/10.1038/s41467-020-14531-6>, 2020.
- 710 Tuyls, K. and Parsons, S.: What evolutionary game theory tells us about multiagent learning, *Artificial Intelligence*, 171, 406–416,
<https://doi.org/10.1016/j.artint.2007.01.004>, 2007.
- United Nations: United Nations Framework Convention on Climate Change, fCCC/INFORMAL/84, GE.05-62220 (E) 200705, 1992.
- United Nations: Paris Agreement, adopted 12 December 2015, FCCC/CP/2015/L.9/Rev.1, 2015.
- United Nations Environment Programme: Emissions Gap Report 2024: No more hot air ...please!, Tech. rep., UNEP, Nairobi,
715 <https://doi.org/10.59117/20.500.11822/46404>, 2024.
- Unruh, G. C.: Understanding carbon lock-in, *Energy Policy*, 28, 817–830, [https://doi.org/10.1016/S0301-4215\(00\)00070-7](https://doi.org/10.1016/S0301-4215(00)00070-7), 2000.



- van der Vaart, E., Beaumont, M. A., Johnston, A. S. A., and Sibly, R. M.: Calibration and evaluation of individual-based models using Approximate Bayesian Computation, *Ecological Modelling*, 312, 182–190, <https://doi.org/10.1016/j.ecolmodel.2015.05.020>, 2015.
- 720 Vasconcelos, V. V., Santos, F. C., Pacheco, J. M., and Levin, S. A.: Climate policies under wealth inequality, *Proceedings of the National Academy of Sciences*, 111, 2212–2216, <https://doi.org/10.1073/pnas.1323479111>, 2014.
- Vehtari, A., Gelman, A., and Gabry, J.: Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC, *Statistics and Computing*, 27, 1413–1432, <https://doi.org/10.1007/s11222-016-9696-4>, 2017.
- 725 Weitz, J. S., Eksin, C., Paarporn, K., Brown, S. P., and Ratcliff, W. C.: An oscillating tragedy of the commons in replicator dynamics with game–environment feedback, *Proceedings of the National Academy of Sciences*, 113, E7518–E7525, <https://doi.org/10.1073/pnas.1604096113>, 2016.
- Wunderling, N., Donges, J. F., Kurths, J., and Winkelmann, R.: Interacting tipping elements increase risk of climate domino effects under global warming, *Earth System Dynamics*, 12, 601–619, <https://doi.org/10.5194/esd-12-601-2021>, 2021.
- Yao, Y., Vehtari, A., Simpson, D., and Gelman, A.: Using Stacking to Average Bayesian Predictive Distributions (with Discussion), *Bayesian Analysis*, 13, 917–1007, <https://doi.org/10.1214/17-BA1091>, 2018.