Multi-Agent Planetary Resource Negotiation Engine (PRNE) for Equitable Global Resource Allocation

Raghu Para*

Independent Researcher, Ontario, Canada

Abstract. This paper introduces the Planetary Resource Negotiation Engine (PRNE), a system designed to help both human and non-human agents negotiate the use of Earth's limited resources under strict ecological boundaries. Unlike conventional models that treat environmental limits as secondary concerns, PRNE places planetary thresholds—such as climate, water, and biodiversity—at the core of its decision-making. The framework models ecosystems themselves (forests, rivers, coral reefs) as active agents with the power to veto harmful proposals or impose costs when ecological limits are threatened. PRNE blends market-style trading with treaty-like rules to promote fairness and protect future generations. Instead of relying on economic discounting, the model ensures minimum rights for future utility. Mathematically, the engine is built as a constrained multi-agent Markov decision process, solved using a mix of equilibrium search and constrained reinforcement learning. A prototype scenario involving nations, corporations, NGOs, and ecosystem agents at the water–energy–food nexus is presented to test the design. Results include resource allocation schedules, ecological boundary usage, and equity measures, with stress tests under shocks such as droughts and tipping points. The paper concludes with governance and ethical considerations to ensure PRNE is transparent, fair, and responsibly applied. **Keywords:** Planetary Boundaries, Multi-agent Systems, Climate Negotiation, Mechanism Design, Intergenerational Fairness, Sustainability

1 Introduction

Humanity has already pushed the Earth system beyond its "safe operating space": a 2023 reassessment finds six of nine planetary boundaries—among them climate change, biosphere integrity, and freshwater use—are now transgressed. Atmospheric CO_2 concentrations exceed safe thresholds, and human pressures on the biosphere have crossed high-risk levels [1]. At the same time, billions still lack necessities, producing a dual crisis of ecological overshoot and social shortfall that current institutions struggle to reconcile [2]. Conventional governance tools underperform: Integrated Assessment Models (IAMs) typically optimize economic objectives and only later incorporate environmental limits as soft or secondary

*Corresponding author: rpara@ieee.org

constraints, diluting biophysical realism. Meanwhile, international processes under the UN, despite milestones like the Paris Agreement, have not secured pledges consistent with safe temperature pathways, and feasibility analyses indicate formidable barriers to achieving 1.5–2 °C stabilization under stated policies.

This context has catalyzed calls for fundamentally new approaches to managing global commons under deep uncertainty. Stern and colleagues argue that the economics of climate action must be reframed to reflect tail risks and urgency, while sustainability science emphasizes that human well-being depends on coupling Earth-system resilience with justice. Planetary justice advances the idea of fairly distributing "ecological space" within strict biophysical limits. Aligned with this view, emerging proposals for "safe and just Earth system boundaries" imply equity-aware thresholds that can be stricter than conventional targets—for example, limiting warming to ~1.0 °C to avoid significant harms to tens of millions when justice is explicitly considered [3].

The paper introduces the Planetary Resource Negotiation Engine (PRNE), an integrative simulation framework for allocating finite resources—carbon budgets, freshwater withdrawals, and land use—under binding Earth-system constraints. PRNE enforces planetary boundaries as hard feasibility limits within the optimization and bargaining space (e.g., cumulative CO_2 within the remaining budget; basin-level extraction below sustainable yield; habitat above minimum intact-area thresholds). Within this feasible region, heterogeneous, self-interested agents negotiate across coupled resource issues and make interdependent trade-offs. By encoding biophysical limits ex ante, PRNE aims to expose policy pathways that keep the planet within a "safe and just corridor" while meeting human needs.

PRNE's design [4] rests on four pillars: (a) Ecosystems as Agents—critical subsystems (forests, oceans, watersheds) are given formal agency via utility functions tied to ecological integrity (e.g., carbon storage, biodiversity), with the capacity to veto or penalize proposals that risk irreversible degradation, operationalizing rights-of-nature and stewardship principles; (b) Hard Earth-System Constraints—limits are embedded directly, not imposed ex post through prices or penalties, ensuring resilience by construction; (c) Hybrid Market–Treaty Mechanism—market instruments (auctions, tradable permits, externality pricing) are co-optimized with treaty features (negotiated quotas, side-payments, and non-market values) to balance efficiency, equity, and enforceability; (d) Intergenerational Equity Module—explicit minimum thresholds on future welfare and resource stocks move beyond pure discounting and function as commitment devices shown to strengthen long-term cooperation

2 Background and Related Work

2.1 Climate Economics and Negotiation Models

For decades, Integrated Assessment Models (IAMs) like DICE and RICE have guided climate policy by linking simplified climate dynamics with economic growth models. These classical IAMs generate cost-benefit-optimized pathways—such as optimal carbon price trajectories—under the assumption of a single global planner, omitting the strategic behavior of sovereign actors. Climate action is a game-theoretic challenge: many independent players, no central authority, and strong free-rider incentives [5].

Game-theoretic climate models have explored coalitions and treaties, but often with stylized assumptions. Nordhaus's "Climate Club," combining an internal carbon price with external tariffs for non-members, exemplifies cooperation through self-interest alignment. More recent work integrates multi-agent reinforcement learning (MARL) into IAMs, enabling regional agents to learn negotiation strategies. Zhang et al.'s RICE-N extends RICE

with n strategic agents that negotiate emissions cuts; MARL agents formed clubs and bilateral deals, reducing peak warming with fairer cost-sharing and modest GDP impacts compared to fixed-policy baselines. The introduction of a simple climate-club protocol significantly improved both equity and climate outcomes.

Agent-based models (ABMs) have also been used to represent heterogeneous actors and bounded rationality, producing emergent phenomena beyond top-down optimization. Traditional ABMs, however, often fix agent decision rules and struggle with normative constraints. MARL addresses this by letting agents learn policies via trial-and-error. The 2022–2023 AI for Global Climate Cooperation initiative used MARL in RICE-N, showing that carefully designed reward structures—such as penalizing inequality—can boost cooperation. Still, most MARL climate models rely on soft enforcement via penalties rather than binding feasibility constraints.

2.2 Planetary Boundaries and System Constraints

The Planetary Boundaries (PB) framework defines limits for key Earth system processes—climate, biosphere integrity, freshwater, land use, biogeochemical flows, ocean acidity, aerosol loading, and novel entities—whose breach risks destabilization. As of 2023, six boundaries are exceeded, including GHG concentrations, biodiversity loss, deforestation, nutrient cycle disruption, and altered freshwater flows; ocean acidification is near the threshold.

Researchers now call for embedding PB limits in decision-making at all scales [3] advance an "Earth system justice" framework to fairly share the remaining ecological budget, while others emphasize equitable well-being within biophysical limits. PRNE builds directly on these concepts, operationalizing PBs as explicit nonlinear constraints in agent negotiations. Advances in Earth system modeling now quantify global limits and downscale allocations—for example, translating the global carbon budget into national quotas using equity principles [6].

PRNE's inclusion of ecosystem agents marks a novel step. The Rights of Nature movement, recognized in countries from Ecuador to New Zealand, grants legal standing to rivers and forests [7]. Climate ethics similarly explore guardians for future generations and ecosystems. Researchers also argue for non-human nature as active political participants. In PRNE, ecosystems hold utility functions (e.g., a forest's carbon storage and biodiversity) and can negotiate or enforce conditions. This parallels "Future Design" experiments in Japan, where representing future generations in negotiations led to more sustainable decisions.

2.3 Mechanism Design and Policy Instruments

Ensuring negotiated outcomes are efficient, stable, and fair requires careful mechanism design. Tools like the Vickrey-Clarke-Groves (VCG) mechanism can elicit truthful preferences and efficient allocations with externalities, though issues like budget imbalance and collusion remain. Environmental economists have explored incentive schemes to foster cooperation.

PRNE incorporates VCG-style externality pricing where possible: e.g., a Pigouvian transfer if a country's plan reduces river flow, credited to ecosystem restoration. Where VCG is impractical, budget-balanced alternatives or hybrid cap-and-trade systems are used. PRNE's negotiation protocol is multi-round and multi-issue—agents simultaneously bargain over emissions, land, water, and conservation funding—allowing package deals that can improve agreement feasibility.

Transparency is essential: each PRNE proposal is annotated with metrics on winners/losers, boundary pressures, and compensations. For example: "Proposal X meets all

energy needs but brings Basin Y's freshwater use to 95% of its limit—triggering $Z investment in recycling." Such explainability ensures trust, identifies "efficient but unjust" outcomes, and aligns with principles of sustainable AI that emphasize augmenting human decision-making [8].

3 Proposed Framework and Methodology

3.1.1 System Overview

The Planetary Resource Negotiation Engine (PRNE) (Fig. 1) is a computational multi-agent framework that integrates human-centric agents, ecosystem agents, and facilitator agents to negotiate over finite planetary resources such as emission budgets, freshwater rights, land-use allowances, and fishing quotas. Unlike traditional economic optimization models, PRNE operates under immutable ecological constraints derived from planetary boundaries—covering climate stability, biosphere integrity, land-system change, freshwater use, nitrogen and phosphorus cycles, ocean acidification, aerosol loading, novel entities, and ozone integrity. Negotiations proceed in cycles, beginning with an algorithmic market allocation that maximizes a global welfare objective subject to these hard limits, followed by a treaty-adjustment phase where agents re-allocate resources to address equity concerns through side payments, technology transfers, or compensatory actions. Ecosystem agents enforce "red lines" by rejecting infeasible deals, while future-generation agents ensure minimum long-term reserves—embedding intergenerational justice directly into the decision-making process. This dual-phase approach ensures both efficiency and fairness while preventing ecological overshoot.

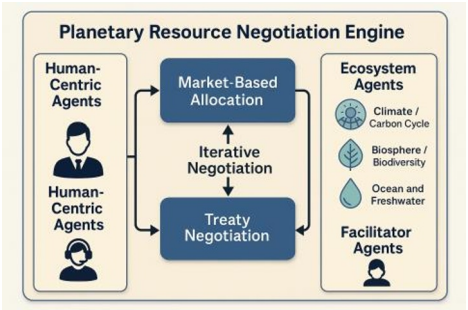

Fig.1. Planetary Resource Negotiation Engine Framework.

3.1.2 Architecture

PRNE's architecture comprises four core components (Fig. 2): the World State Module, which tracks planetary system status, resource stocks, and agent utilities; the Market Solver, an optimization engine (linear, nonlinear, or mixed-integer) that produces initial allocations under strict ecological and intergenerational constraints; the Negotiation Simulator, which allows agents to propose, exchange, and evaluate package deals across multiple issues simultaneously; and the Mediation Agent, an AI-driven facilitator that generates Pareto-improving proposals using bargaining models such as Nash or Kalai–Smorodinsky.

The system operates periodically, enabling adaptive re-negotiation as new data on planetary boundaries or socio-economic conditions emerges. By including ecosystem and future-generation agents alongside human actors, PRNE captures both biophysical realities and ethical imperatives. In practice, negotiation bundles may combine emission quotas, land-use rights, water allocations, and conservation funding into a single deal space, enlarging

opportunities for agreement. The architecture supports hybrid computation—exact solvers for tractable sub-problems and multi-objective evolutionary algorithms or reinforcement learning for complex bargaining scenarios—ensuring scalability from small regional settings to global simulations.

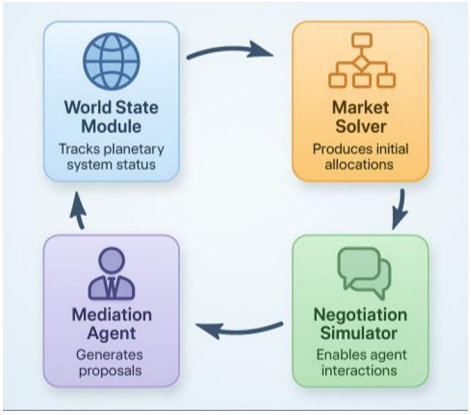

Fig. 2. Interactions of PRNE Components.

3.1.3 Mathematical Formulation

Let the agent set be $\mathcal{N} = \{1, \ldots, N, E_1, \ldots, E_B, F_1, \ldots, F_G\}$, where N denotes human agents, E_k ecosystem agents for B planetary boundaries, and F_h future-generation agents. Each agent i receives an allocation vector $x_i = (x_{i1}, \ldots, x_{iM})$ of M resources, forming allocation matrix X. For each planetary boundary k, the resource pressure function R_k (X) must satisfy $R_k(X) \ll L_k^{max}$, where L_k^{max} is the safe limit. Utilities $U_i(x_i)$ for human agents increase with beneficial resources and decrease with burdens, while future-generation utilities $U_{F_h}(X)$ depend on remaining long-term ecological capacity. A floor constraint $U_{F_h}(X) \gg U_{F_h}^{min}$ ensures intergenerational protection.

The system operates in two stages. Stage 1 (Market Optimization) maximizes a weighted sum of human utilities $\sum_{i=1}^{N} \pi_i U_i(X)$ subject to planetary and intergenerational constraints, yielding an initial allocation X^*.

Stage 2 (Negotiation Adjustment) solves a Nash bargaining problem $\max \pi_{i \in N} [U_i(X) - U_i^d]$ for participating agents N where U_i^d denotes disagreement utilities, again subject to the same hard constraints.

Shadow prices λ_k for each boundary and μ_k for each intergenerational floor quantify the marginal value of relaxing constraints, effectively revealing implicit carbon or water prices. This formulation accommodates both convex optimization for concave utilities and heuristic search for non-convex negotiation spaces, enabling flexible deployment on diverse computational platforms.

4 Experimental Design and Reliability

4.1.1 Resource and Allocations

Let M denote the number of negotiable resource or action variables, such as emission permits, freshwater withdrawal rights, land-use allowances, or cleanup contributions. Each agent i receives an allocation vector $x_i = (x_{i1}, \ldots, x_{iM})$, forming the allocation matrix $X = (x_{ij})$. While some resources may be relevant only to certain agents, still all $x_{ij} \geq 0$ to apply

planetary boundary constraints globally. For each boundary k, a function $R_k(X)$ computes total pressure, bounded above by a safe limit L_k^{max} from Earth system science. Violations render X infeasible, with ecosystem agents E_k enforcing compliance: whether via hard constraints or infinite penalties for overshoot.

Human agents $i \in \{1, \ldots, N, \}$ have utility functions $U_i(X)$ increasing in beneficial allocations and decreasing in burdens, potentially including cross-agent externalities. Future-generation agents F_h derive utility from residual resources or long-term welfare indicators—e.g., the unused carbon budget $L_{CO2}^{max} - R_{CO2}(X)$ or reforestation levels. An intergenerational floor constraint $U_{F_h}(X) \geq U_{F_h}^{min}$ guarantees acceptable future outcomes, such as preserving a fixed share of the carbon budget. While equity constraints among current agents could be added (e.g., per-capita minimums), the present formulation relies on negotiation and transfers to address fairness, focusing instead on binding planetary ceilings and future welfare floors.

4.1.2 Computational Aspects

PRNE can be implemented via constrained optimization or learning-based methods. For concave utilities and convex constraints, centralized solvers such as convex programming or MILP (with discretized allocations) are tractable. For more complex, non-linear cases, multi-objective evolutionary algorithms can approximate the constrained Pareto frontier, from which negotiation solutions are drawn. A decentralized approach could model each party as a reinforcement learning agent proposing or accepting allocation changes, with joint rewards tied to reaching agreement. To manage the vast deal space, negotiation can be restricted (e.g., one resource at a time or proportional adjustments). Initial deployments should target simplified subdomains, scaling gradually as computational capacity and negotiation protocols mature [9].

4.1.3 Key Evaluation Metrics

PRNE is evaluated on: (a) Environmental compliance – strict adherence to all planetary limits, validating enforcement logic; (b) Efficiency – aggregate utility or welfare index relative to optimal feasible benchmarks; (c) Equity – fairness in utility or resource distribution (e.g., Theil index, standard deviation), including intergenerational equity measured against future-agent floors; (d) Stability – game-theoretic robustness where no individual or coalition can deviate under constraints to improve utility [10]. Expected outcomes are Pareto-efficient, constrained-envy-free allocations, where no agent prefers another's allocation when transfers are considered. Stability checks involve verifying that permissible deviations yield no utility gains, indicating equilibrium.

4.1.4 Comparative Baselines

PRNE is benchmarked against: (i) Pure market allocation – efficient but often inequitable or infeasible; (ii) Equal-per-share allocation – equitable but typically inefficient; and (iii) Current policy scenarios (e.g., Paris Agreement pledges, prevailing land-use trends), which studies show will exacerbate boundary overshoots by 2050 [11]. PRNE's negotiated solutions are expected to outperform these baselines by achieving multi-boundary compliance—such as meeting both climate and biodiversity targets—while avoiding the significant overshoots seen in business-as-usual trajectories. Outcomes of PRNE are compared against two simpler allocation mechanisms: a pure market allocation with no negotiation, and an equal-per-share allocation. These baselines highlight the trade-offs between efficiency, equity, and sustainability. Table 1 summarizes how each approach

performs across four key criteria: environmental compliance, efficiency, adaptability to shocks, and transparency/trust.

4.1.5 Performance and Reliability

PRNE must maintain high performance and resilience under uncertainty. Planetary boundaries, agent preferences, and technology costs can shift rapidly, so the system should integrate stochastic or robust optimization to handle worst-case and probabilistic scenarios [12]. For instance, uncertain climate sensitivity could be modeled as a probability distribution, requiring a high likelihood of staying within limits. If economic conditions change post-agreement (e.g., recession), PRNE should reallocate resources without destabilizing the framework.

Table 1. Comparative Outcomes Across Allocation Mechanisms.

Criteria	PRNE	Pure Market Allocation	Equal-per-share Allocation
Environmental Compliance	Within Planetary Limits	Not guaranteed	May meet some limits
Efficieny (Aggregate Utility)	High (Pareto-optimal within constraints)	Highest possible, but ignores equity and limits	Low to moderate
Adaptability to Shocks	High (renegotiation, constraint updates)	Low to Moderate	Moderate
Transparency and Trust	Explainable AI, auditable process	Low transparency	Transparent but simplistic

Technically, it would run as a distributed, always-on platform across multiple nodes for fault tolerance, ingesting large-scale Earth observation data and using high-performance computing for global-scale optimization. Decentralization—where each stakeholder operates its own coordinating agent—reduces single points of failure and mitigates cybersecurity risks while preserving computational efficiency and reliability.

4.1.6 Human-In-The-Loop

PRNE is designed to augment—not replace—human decision-making. It acts as a rapid simulation and analysis engine, enabling negotiators to explore thousands of "what-if" scenarios far faster than traditional talks allow. For example, they could test: If the future generation floor is relaxed from 1.5 °C to 1.6 °C, how much additional GDP is gained now, and is it worth the trade-off? PRNE quantifies such scenarios with Earth system models, shifting debates from rhetoric to evidence-based evaluation. Humans still determine the normative choices—such as equity vs. efficiency priorities—but PRNE ensures they fully grasp the consequences [13]. In effect, it serves as a "flight simulator for Earth," allowing policymakers to navigate complex trade-offs while keeping within planetary boundaries.

4.1.7 Robustness Tests

PRNE can be stress-tested under shocks or uncertainty to evaluate resilience. If planetary boundaries tighten (e.g., stricter scientific limits), ecosystem agents enforce reduced caps, prompting proportional cutbacks or increased compensations. Defection scenarios, where an agent exceeds its allocation, trigger sanctions, renegotiation, or re-optimization among remaining players. Integrated with real-time monitoring (e.g., satellites for deforestation, IoT for emissions), PRNE can detect violations and initiate corrective negotiation or punitive measures. Simulations highlight the role of ecosystem agents in preventing overshoot and the

value of future floors: without them, near-term utility rises at the cost of depleting future resources; with them, long-term welfare is secured with minor present trade-offs [14]. These tests quantify the stability, compliance, and sustainability advantages over unconstrained or business-as-usual approaches.

4.1.8 Data and Monitoring Needs

Deployment of PRNE depends on precise, timely data on resource use and ecological health. Enforcing limits, such as on nitrogen pollution, requires near-real-time monitoring of fertilizer application, runoff, and related indicators. This can be achieved through an integrated Earth observation network combining satellite imagery, atmospheric sensors, and IoT-based water usage tracking. AI analytics would process these streams into actionable insights, feeding a "digital twin" of the planet. PRNE would then translate this dynamic planetary state into informed, constraint-respecting negotiation outcomes.

4.1.9 Transparency and Trust

Building trust in PRNE requires full transparency and accountability in its operations. Stakeholders must be able to verify how allocations are calculated, ensure fairness, and confirm that no manipulation has occurred. This calls for interpretable models or clear post-hoc explanations, such as justifying allocations based on per-capita needs, marginal abatement costs, and adherence to future generation constraints. Open-source governance [15] publishing code and datasets for audit—further enhances credibility. The negotiation process can also be secured through distributed ledger technologies or secure multiparty computation; recording offers and agreements immutably on a blockchain. Smart contracts could then automate enforcement, triggering payments, rewards, or sanctions if commitments are met or breached. Early pilots in carbon trading and fisheries management demonstrate that such mechanisms can provide verifiable, tamper-proof records, bolstering stakeholder confidence in automated planetary governance.

5 Conclusion and Future Directions

This paper presented the Planetary Resource Negotiation Engine (PRNE) as a novel framework for AI-assisted, multi-agent governance of Earth's critical resources under strict ecological constraints. By combining algorithmic market allocation with treaty-style negotiation and embedding planetary boundaries and intergenerational equity directly into the process, PRNE addresses the grand coordination challenge of enabling human development for all while keeping within the safe and just operating space for humanity. It aims to operationalize sustainable multi-agent systems at a planetary scale through rigorous computation and equitable negotiation dynamics.

A key innovation is treating ecosystems and future generations as first-class negotiating agents alongside human stakeholders. PRNE is formalized as a constrained multi-agent optimization problem, encoding planetary boundaries as hard constraints and future welfare floors as non-negotiable minima. The design allows iterative market-based solutions refined through bargaining, producing outcomes that are both efficient and fair. In relation to prior work, PRNE builds on advances in AI for climate cooperation, agent-based sustainability modeling, and planetary boundary science, synthesizing these into a unified decision-support framework.

Our conceptual simulations indicate that PRNE can prevent tragedy-of-the-commons scenarios by enforcing non-negotiable ecological limits and directly counteract tragedy-of-

the-horizon tendencies by safeguarding future interests. By blending market efficiency with mediated negotiation, PRNE can yield Pareto-optimal solutions acceptable to all stakeholders—where traditional political processes often stagnate in zero-sum deadlock.

Future work spans technical, theoretical, and institutional dimensions. Technically, ecosystem agents could be linked to dynamic Earth system models that adjust constraints or utilities based on projected outcomes, effectively creating a feedback loop between real-time simulation and negotiation. This aligns with "digital twin" Earth initiatives, enabling PRNE to test thousands of policy scenarios rapidly. Agents could also incorporate adaptive preferences, reflect societal value shifts, and support the dynamic entry or exit of stakeholders. From a game-theoretic standpoint, stability under coalition formation must be examined: could subsets of agents collude for greater benefit, undermining the grand coalition? Addressing this involves mechanism design for coalition stability, ensuring PRNE outcomes lie within the game-theoretic core through well-calibrated transfers and benefit-sharing.

Institutionally, legitimacy and inclusivity are paramount. PRNE should incorporate non-market voices, such as indigenous communities, local stakeholders, and civil society, via dedicated agents that uphold social floors (e.g., education, health), extending beyond environmental protection to encompass the broader "sustainable development doughnut." This requires normative consensus on minimum social thresholds and methods for measuring them alongside ecological indicators.

Ethical alignment of PRNE's AI mediator is critical. The system's algorithms must be transparent, explainable, and resistant to manipulation, with objective functions aligned to universally agreed goals of sustainability and equity. Governance of PRNE itself warrants careful design; it should function as a global public good, open-source, and overseen by an international, multi-stakeholder body—possibly under a UN mandate or a governance model like the IPCC but with operational authority.

The challenges of the 21st century—climate change, biodiversity loss, inequality, and the need to protect future generations—demand integrated governance systems that transcend the short-termism and fragmentation of historical approaches. PRNE demonstrates how advances in AI and multi-agent systems can help design enforceable, equitable, and adaptive governance mechanisms. By giving nature, a literal seat at the negotiating table and ensuring the future is represented in present decisions, PRNE offers a pathway to align global actions with long-term planetary stability and human well-being.

While significant technical and political hurdles remain, PRNE serves as a computational sandbox for planetary governance design. Its potential lies not just in real-world deployment but in providing a safe, data-rich environment to explore high-stakes trade-offs before they play out. In an era where we are already "well outside the safe operating space for humanity," bold, scalable, and innovative approaches are urgently needed. PRNE, where markets meet treaties under the constraints of hard ecology and intergenerational solidarity, offers one such approach, charting a possible course for a stable and just future on Spaceship Earth.

The challenges of the 21st century: climate change, ecological collapse, inequality, and safeguarding future generations – demand new modes of governance that transcend historical siloed and short-sighted approaches. The paper presents advances in multi-agent systems and AI can be harnessed to create a negotiation engine for the planet, one that rigorously respects the planet's limits as scientifically understood and ensures a fair sharing of the remaining ecological space. By designing systems where nature is not just a constraint but a negotiator at the table, and where the future is represented in today's decisions, we can better align global actions with long-term survival and prosperity. PRNE is a step toward such a system. While many practical hurdles remain before something like it could be implemented in the real world, exploring it now is worthwhile, it provides a computational sandbox for designing the future of global governance. This research sets directions further into interdisciplinary

collaboration between AI researchers, environmental scientists, and policy experts to refine and eventually realize these ideas. The notion of an AI-facilitated planetary negotiation highlights a sense of urgency since we are "well outside of the safe operating space for humanity", innovative solutions at scale are needed. In summary, PRNE represents a novel and promising approach: a multi-agent engine where markets meet treaties under the guidance of hard ecology and solidarity, aiming to chart a stable and just course for Spaceship Earth now and for generations to come.

References

1. Richardson, K., et al. (2023). Earth beyond six of nine planetary boundaries. Science Advances, 9(37). https://doi.org/10.1126/sciadv.adh2458
2. Fanning, A. L., O'Neill, D. W., Hickel, J., & Roux, N. (2021). The social shortfall and ecological overshoot of nations. Nature Sustainability, 5(1). https://doi.org/10.1038/s41893-021-00799-z
3. Gupta, J., et al. (2024). Governance architecture for a safe and just world: Earth system justice respecting safe and just Earth system boundaries while meeting basic needs for all [Report]. Earth Commission.https://earthcommission.org/wp-content/uploads/2024/10/governance_architecture_report_for_publication_v2.pdf
4. Halali, E., & Perez, O. (2025). Binding the future boosts intergenerational sustainability. npj Climate Action, 4(1). https://doi.org/10.1038/s44168-025-00216-7
5. Zhang, T., et al. (2025). AI for global climate cooperation: Modeling global climate negotiations, agreements, and long-term cooperation in RICE-N. OpenReview. https://openreview.net/forum?id=PX29zF9wRb&noteId=hc47FCX7LQ
6. Häyhä, T., Lucas, P. L., van Vuuren, D. P., Cornell, S. E., & Hoff, H. (2016). From planetary boundaries to national fair shares of the global safe operating space—How can the scales be bridged? Global Environmental Change, 40, 60–72. https://doi.org/10.1016/j.gloenvcha.2016.06.008
7. Tănăsescu, M. (2025). The rights of nature go from theory to reality. Current History, 124(858), 21–26. https://doi.org/10.1525/curh.2025.124.858.21
8. Machen, R., & Pearce, W. (2025). Anticipating the challenges of AI in climate governance: An urgent dilemma for democracies. Wiley Interdisciplinary Reviews: Climate Change, 16(2). https://doi.org/10.1002/wcc.70002
9. Donges, J. F., et al. (2020). Earth system modeling with endogenous and dynamic human societies: The copan:CORE open World–Earth modeling framework. Earth System Dynamics, 11(2), 395–413. https://doi.org/10.5194/esd-11-395-2020
10. Sustainability Directory. (2025, May 5). Intergenerational equity policy → Term. Climate Sustainability Directory.https://climate.sustainability-directory.com/term/intergenerational-equity-policy/
11. van Vuuren, D. P., et al. (2025). Exploring pathways for world development within planetary boundaries. Nature.https://doi.org/10.1038/s41586-025-08928-w
12. Motlaghzadeh, K., et al. (2023). A multi-agent decision-making framework for evaluating water and environmental resources management scenarios under climate change. Science of The Total Environment, 864, 161060. https://doi.org/10.1016/j.scitotenv.2022.161060
13. Sustainability Directory. (2025, April 24). Theory of constraints → Term. Sustainability Directory.https://sustainability-directory.com/term/theory-of-constraints/
14. Steffen, W. (2015). Planetary boundaries: Guiding human development on a changing planet. Science, 347(6223). https://www.science.org/doi/10.1126/science.1259855
15. Weidner, T., Galán-Martín, Á., Ryberg, M., & Guillén-Gosálbez, G. (2022). Energy systems modeling and optimization for absolute environmental sustainability: Current landscape and opportunities. Computers & Chemical Engineering, 164, 107883. https://doi.org/10.1016/j.compchemeng.2022.107883