Conservation Strategies in Contested Environments:

Insights from Dynamic Simulations and a Bolivian Case Study*

Diana Weinhold† Lykke E. Andersen‡

January 18, 2026

Draft version – for discussion purposes only

Click here for the latest version of this draft

Abstract

Conservation planning studies typically treat threats as exogenous and evaluate siting rules from a planner's perspective. We argue that conservation is often contested, and develop a sequential land-claim game that models conservation as a dynamic, adversarial contest between conservationists ("Greens") and developers ("Farmers"). We explore the framework in a *Claims World* that isolates the role of rivalry and leakage, and in a *Budget World* that introduces procurement constraints, decomposing outcomes into a Pure Strategy Effect (PSE)—the intrinsic quality of sites a strategy targets—and a Displacement–Leakage Effect (DLE)—the spillover gains from displacing developers' preferred sites when leakage is incomplete. Our results generate several counterintuitive patterns. First, the link between threat-weighting and additionality breaks down once developer adaptation is allowed. Second, reducing leakage can paradoxically increase misallocation. Third, the textbook ratio-greedy rule (maximise efficiency) is systematically *dominated* by the simple value-greedy rule (maximise environment): we explore this 'knapsack reversal' more formally and show how it can produce a 'disappointment gap' between static (Marxan) planning and dynamic implementation. We then transport our dynamic contest to a Bolivia-based planning board constructed from biophysical data and confirm that the qualitative rankings from the simulations carry over, and adversarial outcomes lie well below the static cost-effectiveness upper bound. Tiny-grid equilibria, formal analysis and robustness exercises in the Appendix show that these patterns are consistent with best-response logic rather than artefacts of modelling choices. Together, the results suggest that robust conservation in contested landscapes requires strategies that anticipate adaptation, not just static threats. Replication materials and code are available via GitHub and archived on Zenodo (DOI: 10.5281/zenodo.17114490).

Key words: Conservation strategies, Additionality, Adversarial procurement, Monte Carlo simulation

JEL classifications: Q01, Q57, Q24, C72

*We thank Ben Balmford, Todd Kaplan, Lewis Grant, Timo Goeschl, Ben Groom, and Charlie Palmer for useful comments and suggestions. We thank Fabiana Karina for excellent research assistance. We are grateful to the Natural Environment Research Council (NERC) for funding.

†Dept. of International Development, London School of Economics. Email: `d.weinhold@lse.ac.uk`

‡Universidad Privada Boliviana - SDSN Bolivia. Email: `lykkeandersen@upb.edu`

1 Introduction

Vital ecosystem services and innumerable species are increasingly threatened by global land-use change (Ceballos et al., 2015; Tilman et al., 2017; IPBES, 2019; Almond et al., 2022). In response, more than 190 countries have committed to protecting 30% of the Earth's land and ocean by 2030—the "30×30" goal of the Kunming–Montreal Global Biodiversity Framework agreed in 2022 at the fifteenth Convention on Biological Diversity. Although over 300,000 terrestrial and marine protected areas already safeguard more than 16% of Earth's landmass and 8% of its oceans, meeting the 30×30 target will require nearly doubling terrestrial protection and quadrupling marine protection, demanding an unprecedented acceleration of conservation efforts.

Targeting these new protections efficiently poses substantial financial and logistical challenges, and no consensus has emerged on the most effective siting strategies. Early work emphasized maximizing conservation value (e.g. species representation) while largely ignoring heterogeneity in costs and threats (Margules et al., 1988; Faith, 1992). Later studies incorporated these considerations, with many advocating "hot spot" targeting of high-value, high-risk sites (e.g. Abbitt et al., 2000; Newburn et al., 2006; Game et al., 2008; Venter et al., 2014; Allan et al., 2019; Hansen et al., 2020). Others, however, highlighted that more remote sites of lower ecological quality may be cheaper and therefore more cost-effective to protect (Ando et al., 1998a; Naidoo et al., 2006; Armsworth et al., 2006). More recently, a line of work has circled back to biodiversity-first approaches, with Dinerstein et al. (2017, 2020, 2024) arguing that securing the highest-value ecological sites remains the most reliable strategy.

Empirical studies have likewise produced divergent conclusions. Some document that conservation programs in practice follow Hot Spot principles (Balmford et al., 2000; Wünscher and Engel, 2012; Lu et al., 2023), while others find that protected areas are disproportionately established in low-threat, remote regions—consistent with cost-minimizing or "location-biased" siting (Andam et al., 2008; Joppa and Pfaff, 2011; Pfaff et al., 2015; Baldi et al., 2017; Sims, 2014). For example, Dinerstein et al. (2024) show that since 2018 only a small share of newly protected areas has overlapped with rarity hotspots, underscoring the persistence of location bias in global conservation expansion.

In the last two decades, recognition of this tendency toward low-threat siting, combined with the growing integration of conservation finance with carbon markets, shifted attention toward the principle of "additionality"—the requirement that conservation outcomes exceed the "business as usual" (BAU) trajectory of a non-intervention scenario (Wunder, 2015; Delacote et al., 2024). Additionality has become the dominant evaluation criterion for conservation programs (Andam et al., 2008; Joppa and Pfaff, 2011). Recent evidence from voluntary carbon markets suggests that buyers are increasingly rewarding credits that address concerns over weak baselines and "questionable additionality," with newer vintages and more robust methodologies commanding significant price premiums (Forest Trends' Ecosystem Marketplace, 2023). The appeal of requiring credible additionality lies in directing scarce resources toward immediate threats (Engel et al., 2008; Aspelund and Russo, 2023), yet by channeling investments to frontier lands near agricultural or urban expansion, high additionality often implies sharp socio-economic and political conflicts.

Thus, across the literature there is little consensus on the most effective targeting strategy. Moreover, most academic work has modeled conservation either as an optimal allocation problem for a benevolent planner balancing environmental and economic objectives (e.g. Ando et al., 1998a; Delacote et al., 2024), or as a cooperative framework emphasizing agreements, compensation, or tradable rights (e.g. Ferraro and Simpson, 2002; Engel et al., 2008). These perspectives have yielded valuable insights into cost-effectiveness and institutional design, but they assume away the adversarial dynamics that characterize many real-world conservation conflicts. Indeed, implementing cooperative agreements is challenging - dual-objective projects that simultaneously aim to promote equity and human development alongside conservation frequently achieve suboptimal environmental outcomes (Delacote et al., 2014; Amin et al., 2019) - and political (or physical) conflict around conservation is common around the world. From Brazil, where conflicts between agribusiness interests and environmental agencies have produced abrupt swings in Amazon land-use policy depending on which coalition controls Congress and the presidency (e.g. Fearnside, 2017), to Wisconsin, where the legal battle between conservationists and developers reached the state Supreme Court (Knowles-Nelson Stewardship Program, 2024), land use decision often emerge from political arenas where actors with fundamentally different priorities struggle over institutional levers of power.

Bolivia offers a further case in point. From the creation of Sajama National Park in the

1930s to the wave of protected-area expansion following the 1992 Rio Earth Summit, its conservation landscape has reflected a continuous tug-of-war between pro-development interests (agribusiness, extractive industries, infrastructure) and pressures from civil society and international norms. Protected areas that were once located in remote regions are now increasingly exposed by new roads, oil exploration, and agricultural shifts. In Section 5 we highlight how our analytical framework can shed new light on historical trends in Bolivia and on the performance of alternative conservation strategies on a real, contested landscape.

In this adversarial context, a critical question arises: how do widely-practiced conservation strategies perform when developers actively respond to conservation pressure? Do threat-based targeting rules deliver the additionality they promise? Can cost-effectiveness heuristics guide siting when costs signal rival demand? A small strand of prior work has examined conservation in explicitly non-cooperative settings: Angelsen (2001) models forest appropriation as a contest between state and local actors, showing that outcomes hinge on enforcement and developer responses, while Colyvan et al. (2011) formalize a “conservation game” that demonstrates inefficiency at equilibrium without ranking strategies. Both underscore how conflict can drive outcomes far from the social optimum, but neither provides a systematic comparison of the siting rules debated in conservation practice.

Our contribution is to build a computational framework for comparing conservation strategies under dynamic, adversarial conditions. We model land allocation as a sequential contest between conservationists ("Greens") and developers ("Farmers"), where each side follows a fixed strategy and developers may re-target when blocked. We compare several widely-discussed conservation approaches—maximizing environmental value, targeting biodiversity hotspots, blocking high-threat development, and maximizing cost-effectiveness—across varying levels of "leakage" (the degree to which developers can substitute alternative sites when their preferred projects are blocked). Farmers, in turn, can pursue a simple profit-maximizing strategy, or a 'strategic' approach that prioritizes plots that are at higher risk from being claimed by the Greens. Because Farmers also follow fixed strategies, we can calculate the Farmers' preferred 'Business as Usual' (BAU) trajectory of land use, and directly observe deviations from BAU induced by Green strategies. Monte Carlo simulations then assess how the Green strategies perform against both Farmer strategies in terms of conservation outcomes, additionality, and welfare.

We explore this contest under two institutional variants. In Claims World, teams allocate a fixed number of "claims" (one claim per plot), isolating the core mechanics of sequential competition and leakage. In Budget World, teams purchase plots subject to budget constraints, with prices equal to agricultural values. This introduces an adversarial procurement dimension and stress-tests whether Claims World insights survive under explicit resource constraints. In both Worlds the framework isolates two distinct mechanisms: the Pure Strategy Effect (PSE), which reflects the intrinsic quality of sites a strategy directly targets, and the Displacement–Leakage Effect (DLE), which captures additional gains when blocking developers BAU plots under imperfect leakage. This decomposition clarifies why strategies succeed or fail in ways that aggregate outcome measures obscure.

This approach differs from the usual game-theoretic focus on deriving equilibrium strategies in closed form: for the dimensionality of our games, full analytic solutions are intractable. Instead, we run a computational horse race, evaluating how fixed strategies perform when subjected to adversarial pressure and endogenous developer responses. Mechanically, our model is best viewed as a two-player *alternating draft* (picking-sequence) game over indivisible plots (Brams and Kaplan, 2004; Budish and Cantillon, 2012). In *Claims World*, each side holds a fixed stock of claims (one claim secures one plot) and the teams take turns selecting remaining plots until their claims are exhausted. In *Budget World*, the same alternating-pick logic is combined with procurement constraints: teams purchase plots subject to remaining budgets, with prices tied to developers' valuations. Leakage makes this draft-like contest state dependent: when Greens preempt plots that Farmers would have chosen under their business-as-usual plan, Farmers' effective future choice set is reduced via claim/budget deductions when leakage is incomplete. This sequential-draft structure captures the core tension between offensive play (securing intrinsically valuable sites) and defensive play (preempting rival targets), while allowing the opportunity set to evolve endogenously over play.

Our simulations yield several striking and robust findings. In Claims World, threat-weighting does not reliably produce additionality once developer re-targeting is endogenised. Under high leakage, pure threat-chasing (blocking developers' preferred sites) performs no better than random protection, and can even do worse when it diverts effort away from intrinsically valuable sites. Moreover, reducing leakage can paradoxically increase welfare losses by shifting high-agricultural-value plots into conservation without commensurate ecological gain. Because developers can no longer substitute freely, successful blocking increasingly implies converting plots that optimally belong in agriculture, even when the environmental gains are modest.

Budget World reveals perhaps our most striking result - we uncover a dynamic knapsack reversal: once developer re-targeting and adversarial play are modeled explicitly, a simple value-first rule ('buy the highest-E sites') systematically outperforms the canonical ratio-greedy cost-effectiveness heuristic often used in much of the planning literature as a simplified analogue of static Marxan-style planning under a passive-developer assumption. When cost signals where rivals are concentrating, contesting expensive sites is more effective than avoiding them. This pattern is conceptually related to insights from contest theory and bilevel knapsack models, where competition is endogenously drawn to attractive opportunities and 'efficient' targets can become overcontested (Kovenock and Roberson, 2010; Pferschy et al., 2021; Fischetti et al., 2019). However, to our knowledge it has not previously been demonstrated in a two-sided sequential contest, nor in the context of conservation planning heuristics.

To ground these computational findings, we connect simulation patterns to best-response logic in tractable polar cases (Appendix A8). At full leakage ($L = 1$), optimal Green play reduces to a myopic preempt-or-value rule; at zero leakage (L=0), it becomes an e-first threshold with a BAU bonus. These limiting results rationalize why threat-chasing erodes as leakage rises and why value-first rules dominate against strategic developers. We formalize the Budget World dominance result as Theorem 4.1 (with proof sketch in Appendix A8.9), establishing conditions under which value-greedy purchasing strictly outperforms ratio-greedy selection. Tiny Grid equilibrium exercises (Appendix A8.3) demonstrate these are not artifacts of modelling choices or Farmer heuristics, but reflect robust features of adversarial sequential contests.

Finally, to bring our stylized computational framework to real data, we then exploit our rich high-resolution maps of ecosystem-service values and agricultural potentials in Bolivia to construct a hydrologically coherent planning-unit (PU) "board" for the country, compute PU-level environmental and agricultural scores, and re-run the competition between heuristic Green strategies facing both naïve and strategic Farmers. This approach also allows us to further compare our dynamic Budget World outcomes to a static Marxan-style cost-effectiveness benchmark and demonstrate that the patterns observed in stylised Monte Carlo grids—most notably the dominance of value-first over ratio-greedy rules and the fragility of pure threat-chasing—persist when the contest is played out on a landscape derived from real bio-physical patterns.

The remainder of the paper proceeds as follows. Section 2 sets up the baseline *Claims World* contest, defines leakage, and introduces the Green and Farmer heuristics together with the PSE/DLE accounting. Section 3 reports the baseline Monte Carlo results across leakage levels and developer behaviours, covering conservation, additionality, and welfare trajectories. Section 4 extends the contest to a *Budget World* with prices equal to agricultural values, states and discusses Theorem 4.1 on leakage-invariant dominance in Purchased Conservation, and compares Max Environment to Max Efficiency. Section 5 then turns to Bolivia. Section 5.1 briefly summarises the evolution of the Bolivian protected-area network and its interaction with agricultural expansion, while Section 5.2 uses the same data to build a Bolivian planning board, derive PU-level environmental and agricultural scores, run the Green strategy competitions, and compare dynamic Budget World portfolios to a static Marxan-style cost-effectiveness plan. These exercises show that the qualitative rankings from the simulations carry over to the Bolivia landscape, and they quantify how far adversarial outcomes sit below the static upper bound. Section 6 concludes and discusses policy implications and broader lessons for conservation planning under adversarial dynamics.

Robustness and extensions (correlated values, unequal power, heavy tails, and spatial externalities) appear in Appendices A2–A6; the formal Budget World setup and theoretical results are in Appendices A7–A8; Appendix A9 supplements the Bolivian case study with descriptive statistics and provides additional details of the construction of the Bolivia Board. Replication materials and links are in Appendix A10.

2 Simulation Framework: Baseline Claims World

In our baseline simulation we explore the outcomes from decisions to conserve or develop land that are inherently dynamic, iterative, and politically contested, modelled as adversarial competition between conservationists and developers. To capture this complexity, we propose a Monte Carlo simulation framework that models the behaviour of two teams: the Greens (conservationists) and the Farmers (developers) who each follow a fixed strategy. We then compare the land use patterns and welfare outcomes resulting from different combinations of Green and Farmer strategy under different levels of leakage.

The Grid

We begin by randomly populating a grid of land “plots” (grid cells) with two attributes: a conservation value and an agricultural value. In the game these values are interpreted as the net present social benefit of conserving or developing a plot, respectively, but more generally they can represent any value space one believes the two teams are contesting. We assume that while the intensity of team preferences may differ from the social valuations, the rank ordering of plots is preserved—farmers prioritize plots in order of agricultural value, conservationists in order of conservation value. The grid serves as a pragmatic abstraction: it contains no spatial information, so cells that are distant on the grid may correspond to adjacent plots in reality, and vice versa. The grid can be initialized to simulate a range of ecological and geographic scenarios by varying the correlation between conservation and agricultural values. In our baseline analysis we assume a correlation of zero, but Appendix Sections A2 and A3 explore outcomes under positive ($\rho = 0.3$) and negative ($\rho = -0.3$) correlation, respectively.

Both plot values can be viewed as comprehensive measures that embed dynamic externalities and other intrinsic characteristics affecting their suitability for conservation or development, and we assume the same rank ordering for private values. For example, while global protected areas have been shown to reduce forest loss on average (Yang et al., 2021), their effectiveness is highly heterogeneous (Leverington et al., 2010; Joppa and Pfaff, 2011; Geldmann et al., 2019; Duncanson et al., 2023; Delacote et al., 2024). In our framework, such heterogeneity is incorporated directly into the conservation value assigned to each plot. A negative correlation between agricultural and environmental values can therefore represent contexts where it is legally or institutionally more difficult to secure conservation outcomes on land with high development potential.

2.1 Claims and Rounds

In Claims World, after the grid is populated, a fixed number of claims—equal to the number of grid cells—are divided between the two teams. One claim secures one plot and in each round the teams then take turns spending their claims to secure plots. Mechanically, then, Claims World is an alternating draft (picking sequence) over indivisible plots: teams take turns selecting remaining plots until claims are exhausted. By varying the number of rounds (for a grid size of 100 the maximum number of rounds is 50), we can adjust the granularity of the game. In our baseline we allow for 50 rounds in which each team spends at most one claim per round. However, it is possible to allow each team to claim a different number of plots per round by varying the allocation of claims and the number of rounds, thereby simulating relative differences in economic or political power. For example in the Appendix section A4 we explore outcomes under an unequal division of power in which Farmers are initially allocated 70% of the claims. We define the Farmer's "business-as-usual" "BAU set" to be the static benchmark or reference point set of plots that Farmers would claim given the Farmer's unconstrained preferences before the game starts. Since the Farmer BAU set is perfectly known at the beginning of the game, it can be used to calculate additionality and leakage, as explained below.

Leakage

Synthesizing evidence across land-intensive interventions, Searchinger et al. (2018) estimate displacement (or 'leakage') rates in the range of 20–50% of conserved area toward alternative locations when land is removed from the feasible development set. Complementing this global perspective, spatial panel estimates for reforestation in the Brazilian Amazon find that approximately 10-15% of reforested area is offset by displaced deforestation in surrounding regions (Silva and Nunes, 2025).

To explore the sensitivity of conservation outcomes to different extents of leakage, we define leakage in the simulation as the extent to which Farmers are able to re-target their development activity when a preferred "business-as-usual" (BAU) plot is pre-emptively claimed by the Greens. At one extreme, 100% leakage implies that Farmers can fully substitute toward alternative plots—either within or outside their original BAU set—so that conservation in one location is offset by development elsewhere. At the other extreme, 0% leakage implies that blocking a BAU plot permanently reduces the total number of plots that can ultimately be allocated to agriculture.

We functionally introduce alternative leakage scenarios by reducing the Farmers' remaining claims at varying rates when the Greens claim plots that the Farmers would have targeted under their preferred BAU land-use pattern. Thus, when leakage levels are 100% there are no deductions and Farmers can spend all their initial claims; if a plot that the Farmers would prefer is captured by the Greens, the Farmers simply claim an alternative plot. On the other end of the continuum, with 0% leakage, one Farmer BAU plot lost to the Greens translates to a one-claim deduction of any remaining unspent Farmer claims. In between, for example with 25% leakage, every four BAU plots claimed by the Greens result in a loss of one unspent Farmer claim. We examine leakage levels of 100%, 75%, 50%, 25%, and 0%.

Game Play

In our simulations, each team may adopt one of several strategies that are outlined below in sections 2.2.2 and 2.2.3. Farmer strategies include Naïve Profit Maximiser and Strategic Profit Maximiser and their strategy determines a fixed preferred BAU set of plots that is determined at the start of play. Green strategies include Max Environment, Block Farmers (Maximize Additionality), Hot Spot, Maximize Difference, and Random. On each grid, we simulate every possible combination of Green and Farmer strategies across different levels of leakage. This experiment is repeated 500 times with a new randomly generated grid in each repetition.

Game Conclusion

The game concludes when both teams have exhausted their claims. If leakage<100% any remaining unclaimed plots are automatically assigned to the Greens in the last round and we then compute several outcomes: (a) the final total conservation score (the sum of environmental values for plots held by the Greens); (b) the final degree of additionality achieved by the Greens (the difference between their actual conservation score and the score they would have obtained had the Farmers claimed all their BAU plots); and (c) the final percentage total welfare loss (from the deviation of the final land use pattern from the social welfare–maximizing allocation). In order to visually differentiate the mechanisms through which strategy and leakage interact we report both the final average outcome scores for every strategy and leakage combination and the dynamic trajectories that the outcomes follow throughout the game.

The assignment of unclaimed plots to Conservation only in the last round is done not because this timing is 'realistic' but rather because we want to clearly distinguish between conservation gains from the strategy and leakage combination itself (what we call the *Pure*

Strategy effect, defined below), and conservation gains from the residual left-over plots (what we call the *Displacement-Leakage* effect, also defined below), which are determined by different mechanisms. More specifically, at different levels of leakage the set of available plots to choose from at any given point changes, so different strategies can lead to different dynamic trajectories. By allocating all the residual unclaimed plots in the last round we can visually separate this more subtle *Pure Strategy* mechanism from the more mechanical *Displacement-Leakage* increase in final Conservation score from residual unclaimed plots, which may also be different across different strategies. Thus, we credit DLE at the end purely as an accounting device so that PSE and DLE are visually separable. Final totals are invariant to this choice. Additionality exhibits no end-round jump because residual plots are almost never in the Farmer BAU set.

2.2 Formal Setup (Baseline Claims World)

The baseline game, which we now describe more formally, is mechanically an alternating draft (picking sequence) over indivisible plots, where teams take turns selecting remaining plots until claims are exhausted. We first generate a 10×10 grid of plots, each with a randomly assigned environmental value and an agricultural value.

Let P represent the set of land plots, where $P = \{p_i\}$, with $N = 100$ plots indexed by $i = 1, \ldots, N$. Each plot p_i is associated with two values:

$$\begin{aligned} e_i &: \text{ the environmental value of plot } p_i, \text{ where } e_i \geq 0 \\ a_i &: \text{ the agricultural value of plot } p_i, \text{ where } a_i \geq 0 \end{aligned}$$

The two land attributes, e_i and a_i, are initialized using a bivariate random number generation process designed to create two values whose correlation can be determined by the user to simulate various underlying ecosystems. Specifically, each pair of land attributes (e_i, a_i) is drawn from a Gaussian copula:

$$(z_e, z_a) \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}\right)$$

where z_e and z_a are standard normal variables, and ρ denotes the latent correlation parameter in the bivariate normal used to generate the Gaussian copula. These are then transformed into uniform variables: $u_e = \Phi(z_e)$ and $u_a = \Phi(z_a)$, where $\Phi(\cdot)$ denotes the cumulative distribution function (CDF) of the standard normal distribution, mapping each z to $u \in [0, 1]$. Finally, u_e and u_a are linearly rescaled to the interval $[0.1, 10.0]$ to produce:

$$e_i = 0.1 + 9.9 \cdot u_e$$

$$a_i = 0.1 + 9.9 \cdot u_a$$

This procedure ensures that the marginal distributions of e_i and a_i are uniform on $[0.1, 10.0]$ and induces dependence through a Gaussian copula with latent correlation parameter ρ. Because the transformations are monotone, the resulting rank dependence is a deterministic function of ρ; in particular, the implied Spearman rank correlation is

$$\rho_S = \frac{6}{\pi} \arcsin\left(\frac{\rho}{2}\right)$$

(and Kendall's $\tau = \frac{2}{\pi} \arcsin(\rho)$). Appendix A5 shows that our results are robust to heavy–tailed environmental draws.

It is important to note that the choice of a grid structure is purely practical — there is no spatial information in the grid. The framework can map to any geographical arrangement or size distribution of plots. We consider the values to be the net present value, including any dynamic externalities, if the plot were protected (for the environmental value) or developed (for the agricultural value), and we assume that the rank ordering of values also encodes the multidimensional concerns of each team that drive their siting preferences.

Guided by the observed sequential evolution of agricultural land and protected areas (for example depicted in Figure 1), we simulate the dynamic decisions of two 'teams', the Farmers and the Greens, competing over a set number of periods to claim plots of land for agriculture or conservation.

Let T_f and T_g represent the two teams, Farmers and Greens, respectively. Each team $T \in \{T_f, T_g\}$ has an initial allocation of claims C_T, where: $C_f + C_g = 100$, i.e. the total number of plots. By varying the allocation of claims between the Greens and Farmers, we simulate differing degrees of political power.

The simulation runs for R rounds. In each round r, Farmers (team T_f) move first, followed by Greens (team T_g). The "Farmers move first" convention is the more conservative test of the value-greedy Green strategy *Max Environment*; tiny-grid tests (Appendix A8.3) illustrate that letting Greens move first gives it a slight advantage by allowing it to secure the top-value plot unopposed, though this advantage dissipates rapidly as grid size increases.

Let C_T^r denote the number of claims team T can use in round r. If the game is played over R rounds, then:

$$C_T^r = C_T/R$$

where C_T is the total number of claims allocated to team T. For example, if $R = 50$, both teams claim 1 plot per round, and if $R = 25$, they claim 2 plots per round.

Let $S_f^r \subset P$ and $S_g^r \subset P$ denote the sets of plots claimed by Farmers and Greens in round r, respectively. The total set of plots claimed by each team after round r is:

$$S_f^{\leq r} = \bigcup_{k=1}^{r} S_f^k \quad \text{and} \quad S_g^{\leq r} = \bigcup_{k=1}^{r} S_g^k$$

A plot p_i can only be claimed once, so $S_f^{\leq r} \cap S_g^{\leq r} = \emptyset \;\; \forall r$.

2.2.1 Leakage

Let S_f^{BAU} represent the set of plots that the Farmers would claim under the 'business as usual' (BAU) scenario if they were allowed to use all their claims without any interference from the Greens. During the game, if the Greens claim a plot that is part of the Farmers' BAU set S_f^{BAU}, then the Farmers must deviate from their business-as-usual desired trajectory. If conservation efforts face 100% leakage, then the Farmers can simply claim an alternative plot elsewhere in the grid that may or may not be in S_f^{BAU}. Thus in the baseline case when there is 100% leakage, both teams spend all their initial claims, resulting in all plots $P = \{p_i\}$ being allocated by the end of the game.

We then vary the degree of leakage, denoted as Leakage $\in [0, 1]$ by reducing the Farmers' remaining unspent claims proportionally for every plot $p_i \in S_f^{BAU}$ claimed by Greens, such that for example:

- For Leakage = 0.5, Farmers lose 1 unspent claim for every 2 plots in S_f^{BAU} claimed by Greens.
- For Leakage = 0.0, Farmers lose 1 unspent claim for every 1 plot in S_f^{BAU} claimed by Greens.

The game concludes when both teams have exhausted their claims, or when no further plots are eligible for claim. If Leakage < 100% and the Greens have claimed plots in the Farmers' BAU set, there may be some residual unclaimed plots at the end of the game, $p_i \notin (S_f^R \cup S_g^R)$ and these are automatically assigned to the Greens in the final round.

2.2.2 Farmer Strategies

Farmers choose plots following two possible simulated strategies:

1. **Naïve Profit Maximizer:** Farmers claim plots based solely on the highest agricultural value, ignoring the environmental value:

$$S_f^r = \max(a_i).$$

2. **Strategic Profit Maximizer:** Farmers recognize that the Greens may want to claim good agricultural plots if they are also environmentally valuable, so they may choose plots with slightly lower agricultural value but higher environmental value to pre-empt likely Green targets. Formally, let $\tau \in (0,1)$ denote the anticipated Green claim share (baseline $\tau = C_g/N$). Farmers rank all plots by environmental value in descending order and treat the top $\lceil \tau N \rceil$ plots—those most likely to be chosen by a value-greedy Green rule—as the "risky" set $\mathcal{R}_\tau$. In each round they claim the available plot with the highest a_i within $\mathcal{R}_\tau$; once $\mathcal{R}_\tau$ is exhausted they continue by descending a_i in the remaining "safe" set $\mathcal{R}_\tau^c$.

$$S_f^r = \begin{cases} \max(a_i^{risky}) & \text{if } |a_i^{risky}| \neq 0, \\ \max(a_i^{safe}) & \text{if } |a_i^{risky}| = 0. \end{cases}$$

This stylized behavior has clear parallels in empirical settings where landowners anticipate conservation interventions and adjust accordingly. For example, studies of the U.S.

Endangered Species Act document preemptive timber harvests or habitat destruction to avoid regulatory restrictions (Lueck and Michael, 2003; Langpap and Wu, 2017).

2.2.3 Conservation Strategies

In each round, Greens evaluate each available plot and choose the plot that maximizes a weighted function of e_i and a_i:

$$S_g^r = \max f(e_i^\alpha, a_i^\beta)$$

where coefficients α and β are respective weights of the conservation and agricultural scores that correspond to different alternative strategies as described below.

1. **Max Environment (value-greedy strategy)** Conservationists claim plots with the highest environmental scores. This corresponds to: $\alpha = 1$ and $\beta = 0$:

$$S_g^r = \max(e_i).$$

2. **Block Farmers (threat-greedy strategy)** Conservationists claim plots with the highest agricultural scores to block farmers from claiming them. This corresponds to $\alpha = 0$ and $\beta = 1$:

$$S_g^r = \max(a_i).$$

3. **Hot Spot (product-greedy strategy)** Conservationists claim plots with both high environmental scores *and* high agricultural scores. The default is $\alpha = 1$ and $\beta = 1$:

$$S_g^r = \max(e_i \times a_i).$$

4. **Maximize Difference (net social benefit strategy)** Conservationists claim plots where the difference between environmental and agricultural values is greatest:

$$S_g^r = \max(e_i - a_i).$$

5. **Random Protection:** Conservationists claim plots at random:

$$S_g^r = \text{rand}(e_i).$$

2.2.4 Baseline Outcomes

We then conduct a comparative statics analysis by running Monte Carlo simulations in which a grid is populated with environmental and agricultural scores according to the specified parameters. Each farmer strategy is run against all Green strategies on the same grid, ensuring comparability. This is repeated 500 times, with the grid randomly repopulated each time, and the average of the following outcomes is recorded.

1. **Final Conservation Score.** The final conservation score is the sum of the environmental values in the plots claimed by the Greens during the game, plus any remaining unclaimed plots at the end that are left undeveloped:

$$Conservation\ Score_{Green} = \sum_{p_i \in S_g^{\leq R}} e_i + \sum_{p_i \notin (S_g^R \cup S_f^R)} e_i.$$

Let $S_f^{\text{BAU}} \subset P$ denote the Farmer BAU set defined on the initial grid. Given this BAU pattern, let $S_g^{\text{BAU}} \subset P$ be the set of plots the Greens would hold if Farmers were allowed to claim all plots in S_f^{BAU}without interference. Final conservation can then be decomposed as the sum of the *Pure Strategy Effect* (PSE) and the *Displacement–Leakage Effect* (DLE),

$$C_{\text{final}} = PSE + DLE.$$

(a) **Pure Strategy Effect (PSE)** — the inherent quality of the conservation targets selected by the strategy:

$$PSE = \sum_{p_i \in S_g^{\leq R}} e_i.$$

(b) **Displacement–Leakage Effect (DLE)** — the additional conservation benefit realized due to effective targeting of Farmer BAU plots under incomplete leakage:

$$DLE = \sum_{p_i \notin \left(S_g^{\leq R} \cup S_f^{\leq R}\right)} e_i.$$

Residual plots are credited at the end of the game purely as an accounting device to separate strategy-driven dynamics from mechanical residual gains. Because residual plots are, for all practical purposes, almost never part of S_f^{BAU}, this convention produces an end-of-game "jump" in *conservation* (through DLE) but not in *additionality.* Final totals are invariant to when these residuals are credited; the convention simply clarifies how the two mechanisms operate visually in Figures 1–6.

2. **Additionality.** Additionality is measured as the actual conservation score achieved by the Greens minus the conservation score they would have obtained had the Farmers been allowed to claim all their BAU plots:

$$Conservation_{BAU} = \sum_{p_i \in S_g^{BAU}} e_i,$$

$$Conservation_{Actual} = \sum_{p_i \in S_g^{\leq R}} e_i + \sum_{p_i \notin (S_g^R \cup S_f^R)} e_i,$$

$$\text{Additionality} = Conservation_{Actual} - Conservation_{BAU}.$$

3. **Total Welfare Loss.** The total welfare loss (%) is the percentage decline from the maximum total welfare that could be achieved to the actual welfare realized:

$$W_{max} = \sum_{p_i \in P} \max(e_i, a_i), \quad W_{actual} = \sum_{p_i \in S_f} a_i + \sum_{p_i \in S_g} e_i,$$

$$\text{Welfare Loss(\%)} = 100 \cdot \frac{W_{max} - W_{actual}}{W_{max}}.$$

Throughout, we use the term *welfare* to denote the total value of plots assigned to their preferred use, measured by $\sum_i \max(e_i, a_i)$ at the social optimum and by realized assignments in the simulations. To the extent that e_i and a_i represent full social values, any misallocation corresponds to a deviation from maximum social welfare. More generally, welfare loss can be read as an efficiency loss within whatever value space one believes the teams are contesting.

2.3 On Heuristics vs. Equilibrium.

Our objective is to conduct a computational comparison of widely discussed conservation strategies and to provide a theoretical interpretation of the results. To establish generality and aid interpretation, Appendix A8 characterizes best-response logic in tractable limiting cases and compares the heuristic strategies with exact equilibria on tiny-grid games. These analyses rationalize and validate the strategy rankings reported in Sections 3 and 4—for example, the erosion of threat-chasing as $L \uparrow 1$ and the dominance of value-first rules against strategic developers—showing that the heuristic comparisons capture essential features of sequential adversarial dynamics rather than artifacts of the computational framework.

3 Baseline Claims World Simulation Results

Our baseline Monte Carlo simulation results compare all Green–Farmer strategy combinations across leakage levels when claims are initially split evenly and $\mathrm{corr}(e, a) = 0$. Appendix A10 provides an interactive browser version of the game; Appendix A2 reports the positive-correlation case ($\rho = 0.3$); Appendix A3 reports the negative-correlation case (weak legal enforcement); and Appendix A4 reports the unequal-power case (more Farmer claims).

3.1 Conservation outcomes of Green strategies

We first examine the total conservation achieved (sum of environmental values in conserved plots). The figures show both final outcomes as well as the dynamic trajectories and Table A1.1 in Appendix A1 reports exact final conservation scores, decomposed into the *Pure Strategy Effect* (PSE) and the *Displacement–Leakage Effect* (DLE), for the three primary Green strategies against the two Farmer strategies at leakage levels of 0%, 50%, and 100%.

Naïve Farmers

Figure 1(a) shows final scores for all five Green strategies (with shading to indicate 95% confidence intervals) and Figure 1(b) plot dynamic trajectories for the three primary strategies against Naïve Farmers.

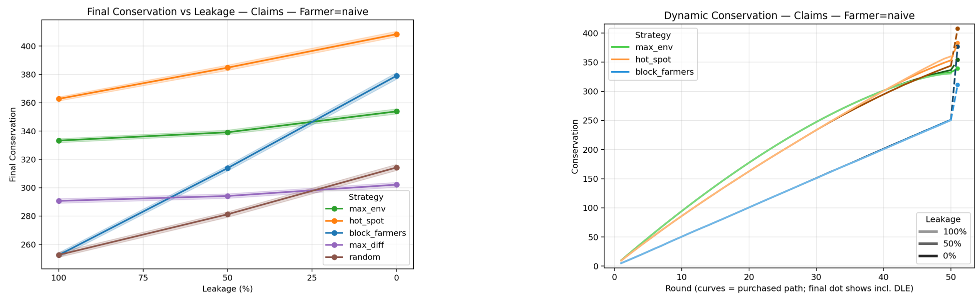

(a) Final Conservation Scores

(b) Dynamic Trajectories

Figure 1: Claims World: Green Strategies and Naïve Farmers — Conservation Scores

As leakage falls, Farmers lose the ability to replace BAU plots captured by Greens, and residual unclaimed plots at the end are assigned to conservation; final conservation therefore rises (though additionality need not, since residual plots are rarely in the BAU set; see Section 3.2). In Figure 1(b), this appears as strategy-specific trajectories (PSE) with differing last-round "jumps" (DLE).

Against Naïve Farmers, *Hot Spot* attains the highest final conservation at all leakage levels; *Max Environment* is second except at the very lowest leakage. The trajectories clarify why: *Max Environment* delivers higher PSE early by focusing solely on environmental value. *Hot Spot* initially sacrifices some PSE to block development on high-agricultural/high-environmental plots, but later reaps gains by securing sites otherwise lost to development; by the final rounds, its PSE surpasses *Max Environment*, with DLE adding a smaller extra margin.

A subtle but informative pattern is that *Hot Spot*'s PSE is lower under 0% leakage than under higher leakage. Because *Hot Spot* weights agriculture and environment equally, greater earlier leakage allows Farmers to remove higher-agricultural plots from the choice set, leaving a later pool with higher average environmental value; under zero leakage, the remaining high-agricultural plots can draw *Hot Spot* away from better environmental options, yielding a nonlinear PSE response to leakage. Even so, *Hot Spot*'s *final* conservation is highest at lower leakage because DLE is larger when more plots remain unclaimed at the end.

By contrast, *Block Farmers* (pure threat-targeting) yields the lowest final conservation for all but the lowest leakage. At 0% leakage it can exceed *Max Environment* in the *final* score, but this advantage is entirely mechanical (DLE): Farmers cannot substitute for lost BAU plots.

Its PSE is consistently below that of *Max Environment*.

These patterns underscore that ultimate success depends on both components: the feasibility of protecting land Greens deliberately claim (PSE) versus the effectiveness of suppressing developer re-targeting (DLE). The end-of-game "jumps" in Figure 1(b) reflect our accounting convention (residual plots credited at once) rather than a behavioral shift. In realistic settings where achieving zero leakage is difficult, PSE will typically matter more than DLE. In the 100% leakage limit, *Block Farmers* performs no better than *Random* when $\mathrm{corr}(e, a) = 0$.

Strategic Farmers

Figure 2(a) shows final scores for all five Green strategies (with shading to indicate 95% confidence intervals) and Figure 2(b) plot dynamic trajectories for the three primary strategies against Strategic Farmers .

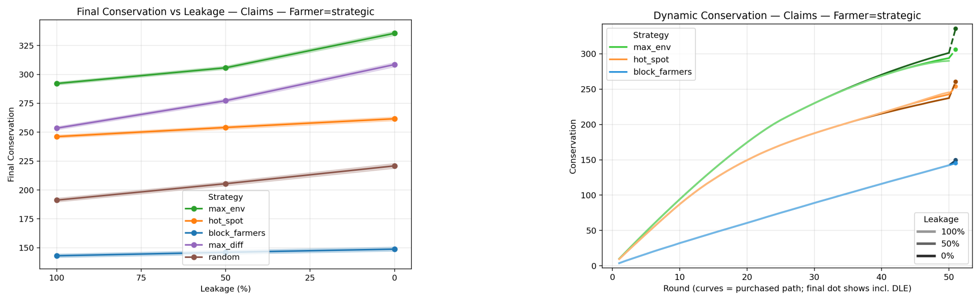

(a) Final Conservation Scores

(b) Dynamic Trajectories

Figure 2: Claims World: Green Strategies and Strategic Farmers — Conservation Scores

Strategic Farmers move some effort from purely highest-agricultural plots toward high-environmental plots that are at risk of being claimed, deferring safer high-agricultural/low-environmental plots to later rounds. With this behavior, the ranking reverses: *Max Environment* strictly dominates in final conservation across leakage levels (Table A1.1; Figure 2(a)). *Block Farmers* is strictly dominated by all other strategies (including *Random*). Figure 2(b) shows the mechanism: when Farmers also target many high-environmental plots, any deviation by Greens from an environment-first rule risks losing those to development. *Max Environment* thus achieves the strongest PSE and, because the Farmer BAU set contains more environmentally valuable plots, a larger DLE than alternatives.

3.2 Additionality of Green strategies

We next consider *additionality*: the realized conservation minus the conservation Greens would obtain if Farmers could claim all BAU plots.

Naïve Farmers

Figures 3(a) and 3(b) present final additionality for all five Green strategies (by leakage) and trajectories for the three primary strategies playing against Naïve Farmers.

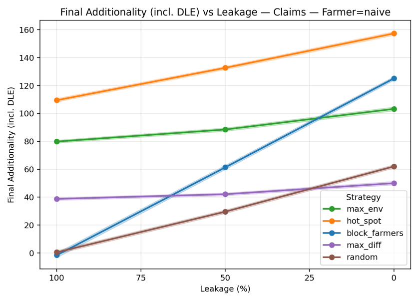

(a) Final Total Additionality

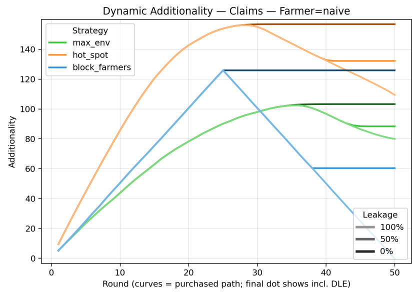

(b) Dynamic Trajectories

Figure 3: Claims World: Green Strategies and Naïve Farmers — Additionality

Final additionality mirrors final conservation because the BAU benchmark is strategy-invariant. Under high leakage, *Block Farmers* achieves essentially no additionality, highlighting that threat-weighting and additionality are not interchangeable. Trajectories are highly nonlinear: additionality only increases when Greens capture a Farmer BAU plot; residual end-of-game plots are almost never BAU, so (unlike conservation) there is no last-round jump.

Leakage is decisive: at higher leakage, later developer re-targeting erodes early gains, and in the 100% leakage limit *Block Farmers* additionality does no better than Random.

Strategic Farmers

Figures 4(a) and 4(b) present final additionality for all five Green strategies (by leakage) and trajectories for the three primary strategies playing against Strategic Farmers.

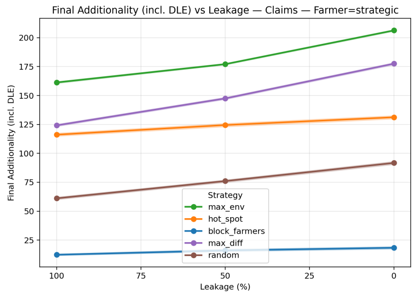

(a) Final Total Additionality

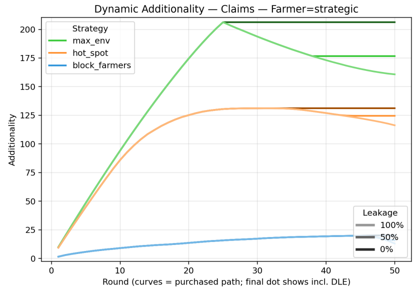

(b) Dynamic Trajectories

Figure 4: Claims World: Green Strategies and Strategic Farmers — Additionality

With Strategic Farmers, levels of additionality are higher than with Naïve Farmers because their BAU sets include more high-environmental plots. *Max Environment* strictly dominates at all leakage levels, followed by *Hot Spot*. *Block Farmers* is strictly dominated by a wide margin. Even though *Block Farmers* and *Hot Spot* are both threat-weighted, they do not necessarily yield greater additionality; indeed, the threat-blind *Max Environment* often secures higher additionality when developers behave strategically.

Synthesis: Claims World Conservation and Additionality

Three general observations emerge. First, despite appearing ex ante to maximize additionality, *Block Farmers* performs poorly: with uncorrelated e and a and iterative play, it often diverts effort to low-environmental plots; under high leakage, Naïve Farmers simply re-target to next-best agricultural plots (sometimes with high e), and Strategic Farmers explicitly chase high-e plots, further undermining measured additionality. Second, the mechanism is *endogenous re-targeting*: once a preferred plot is blocked, developers reallocate, so "threat" is not fixed but adapts to conservation actions. Third, while *Hot Spot* and *Max Environment* are the best performers overall, their relative advantage depends on developer behavior: *Max Environment* dominates against Strategic Farmers (higher PSE and often higher DLE), whereas *Hot Spot* can edge ahead against Naïve Farmers.

Furthermore, in Appendix A8.3 we show that the relative performance of the Green strategies is not an artifact of the specific Farmer heuristics we have chosen. Exact tiny-grid equilibria confirm the ranking and show that in Claims World Hot Spot most closely approximates the equilibrium Green response on larger grids, with Max Environment next best, and Block Farmers collapsing.

3.3 Social Welfare Outcomes of Green Strategies

Finally, we examine overall social welfare outcomes - the shortfall from the maximum attainable total score when each plot is assigned to its highest-value use. Welfare loss occurs from misallocation, whenever Greens claim ('flip') a plot that optimally should be agriculture, or when Farmers flip a plot that should optimally be conserved; the difference between the two values determines the magnitude of the misallocation.

Naïve Farmers

Figures 5(a)–5(b) present the final outcomes and dynamic trajectories of Welfare Loss for Green strategies playing against Naïve Farmers.

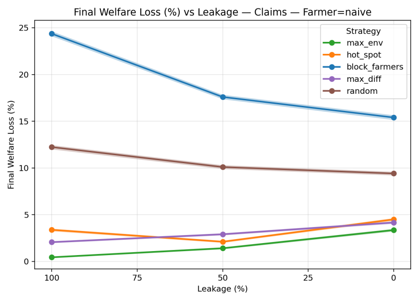

(a) Final Social Welfare Loss (%)

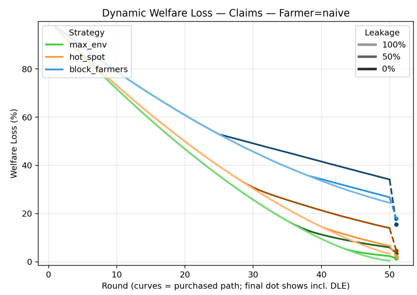

(b) Dynamic Trajectories

Figure 5: Claims World: Green Strategies and Naïve Farmers — Social Welfare Loss (%)

The patterns show that across leakage levels *Max Environment* yields the lowest welfare loss, with *Hot Spot* a close second; *Block Farmers* exhibits much higher losses. Reduced leakage modestly lowers *Block Farmers*' welfare costs (expected, as some unclaimed plots optimally belong in conservation), but produces a counterintuitive pattern for the others: welfare loss rises for *Max Environment* and is U-shaped for *Hot Spot*. The logic is visible in the dynamic trajectories: as the game progresses, strategies that successfully flip more Farmer BAU plots face a greater remaining opportunity set of plots that are disproportionately of high-Agricultural value, increasing the odds that any new Green claim 'flips' a plot and increases welfare loss

(PSE). Then, DLE can further increase losses when residual low-e/high-a plots are forced into conservation at the end. Overall these strategies display greater misallocation at lower levels of leakage.

Strategic Farmers

Figures 6(a)–6(b) present the final outcomes and dynamic trajectories of Welfare Loss for Green strategies playing against Strategic Farmers.

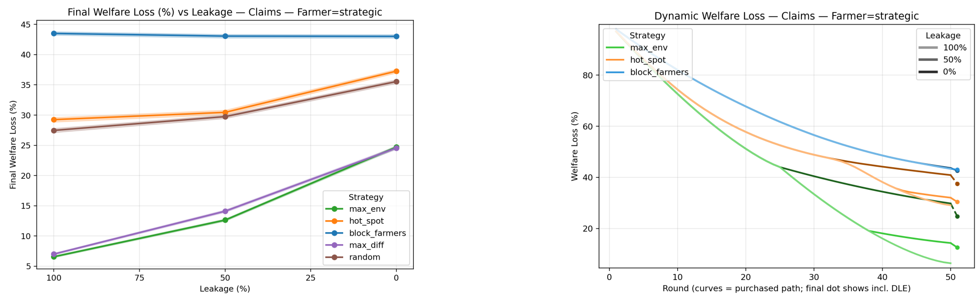

(a) Final Social Welfare Loss (%) **(b)** Dynamic Trajectories

Figure 6: Claims World: Green Strategies and Strategic Farmers — Social Welfare Loss (%)

Note: Shaded bands on static panels indicate 95% bootstrap confidence intervals for the mean across Monte Carlo replications.

Against Strategic Farmers, *Max Environment* again strictly dominates (lowest welfare loss) and *Block Farmers* is worst at all leakage levels. *Hot Spot* exhibits pronounced nonlinearities: early rounds flip plots with small $|a - e|$ (low loss), but as the pool shifts to larger gaps and higher agricultural values, losses accelerate—especially when lower leakage prevents Farmers from reclaiming such plots late in the game. With strategic Farmer behavior the pattern of lower leakage levels raising (via both PSE and DLE channels) late-stage misallocation, and thus welfare loss, is even more striking.

Overall, across both Farmer behaviors, *Max Environment* and *Hot Spot* incur the lowest welfare losses, but paradoxically these losses typically *increase* (welfare worsens) as leakage falls due to misallocation via both shifting opportunity sets (PSE) and higher agriculturally-weighted residual sets (DLE). *Block Farmers* consistently performs poorly, with only limited mitigation as leakage declines.

4 Budget World

In the sections above we have used a stylized, analytically transparent “Claims World” to isolate the role of rivalry, highlighting how leakage, threat-weighting, and strategic behaviour shape conservation outcomes in an adversarial framework. In this section we extend the contest to a “Budget World” in which both Farmers and Greens purchase plots subject to fixed initial budgets and plot-specific costs (agricultural values a_i), preserving the adversarial contest while introducing an alternative, continuous rather than discrete, allocation mechanism for procurement to isolate the informational channel by which prices proxy for rival demand.

The formal set-up of Budget World is outlined in Appendix A7 and follows intuitively the logic developed in Claims World, with one key change: the price of each plot is set equal to the Farmers' valuation, a_i. As in Claims World, in our baseline Budget World analysis environmental and agricultural values are independent draws, so the price faced by Greens is perfectly informative about the Farmers' payoff but, absent correlation, uninformative about Greens' own payoff. This two-space structure—where rival payoffs are orthogonal rather than shared—differs from standard one-space contest and interdiction/knapsack settings. More generally, Budget World studies contests in which each item has a value for the planner and a value for a rival, and prices are (weakly) increasing in the rival's value; our conservation application is a stark special case, with price literally equal to the Farmers' value and insensitive to Greens' valuation except through whatever correlation happens to exist between them.

To run the Budget World competition we calculate the amount required to purchase the whole grid and then split that amount into a Farmer budget B_f^0 and a Green budget B_g^0 (the default split is 50–50). Each team then takes turns purchasing a plot subject to their remaining budget. Each round increments when one or both teams makes a purchase, so the total rounds are endogenous and often exceed 50 because some cycles are one-sided purchases. If Greens purchase a Farmer-BAU plot p_i, Farmer's remaining budget is reduced by a fraction of the price (agricultural value a_i) of the plot, with the fraction given by the level of leakage as in Claims World.

Farmer strategies in the Budget World closely follow those in the Claims World:

1. **Naïve budgeter**: purchase the highest-a_i plot feasible within B_f^0.
2. **Strategic budgeter**: identify "risky" plots, defined as those that would be chosen by a Green Max-E strategy under B_g^0; purchase risky plots in descending a_i, then purchase safe high-a_i plots.

Likewise, Green strategies in the Budget World closely follow those in the Claims World:

1. **Max Environment**: The value-greedy strategy: in each round the Greens purchase the single affordable plot with the highest environmental value e_i (the greedy rule).
2. **Hot Spot**: The product-greedy strategy: the Greens target plots with the highest product $e_i \times a_i$.
3. **Block Farmers**: The threat-greedy strategy: the Greens purchase the affordable plot with the highest agricultural value a_i.

In addition, the move to Budget World allows us to explore a possible fourth strategy:

4. **Max Efficiency**: The ratio-greedy strategy: the Greens purchase the plot with the highest ratio $r_i = e_i/a_i$ (the *ratio-greedy* heuristic solution to a budget-constrained maximization of environmental value).

The accounting logic in Budget World parallels that of Claims World but adapts to budget-based play. We distinguish *Purchased Conservation* (PC)—the analogue of the Pure Strategy Effect (PSE)—from the *Displacement–Leakage Effect* (DLE), which captures the additional conservation value of plots that remain undeveloped when Green purchases reduce Farmers' effective capacity to complete their BAU plans. DLE is reported only under budget-parity ($s = 0.5$), where residual plots represent displaced Farmer opportunities rather than unequal initial resources. Dynamic trajectories for PC therefore show no end-of-game jump, while final *conservation* and, under parity, final *welfare* can exhibit a discrete increase in the last round when residual plots are credited. This convention mirrors the Claims World treatment and ensures that PSE/PC dynamics are visually separable from mechanically induced DLE effects. We calculate *Additionality* on an event basis: it increases by e_i when Greens purchase a Farmer BAU plot and decreases by e_i when Farmers purchase a plot that Greens would have purchased under BAU. Similarly, we distinguish between *Welfare Loss* from purchased plots only and final *Welfare Loss* (in the budget-parity case only) that includes DLE effects.

4.1 Budget World Results and Discussion

Extending the adversarial contest from *Claims World*, the baseline that isolates the role of rivalry, to *Budget World*, which introduces the element of procurement under constraints, confirms the robustness of the central mechanisms and contributes important new insights by allowing us to explore the relative performance of *Max Efficiency* in an adversarial framework.

Budget World: Conservation outcomes

Figures 7–8 demonstrate that the core Claims World findings carry through unchanged. Threat-weighting and conservation do not coincide: Hot Spot strategies deliver higher (purchased) additionality against Naïve farmers by capturing BAU plots, but collapse against strategic farmers, while Max Environment consistently dominates on long-run conservation outcomes. Likewise, strategic behavior by Farmers sharpens relative performance differences. A parallel set of tiny–grid equilibria (Appendix A8.3) for Budget World shows that *Max Environment* increasingly outperforms the ratio–greedy *Max Efficiency* rule as dimensionality grows, confirming that the dominance result reported in our main Monte Carlo simulations is not an artifact of weak Farmer behavior but persists even against optimizing Farmers.

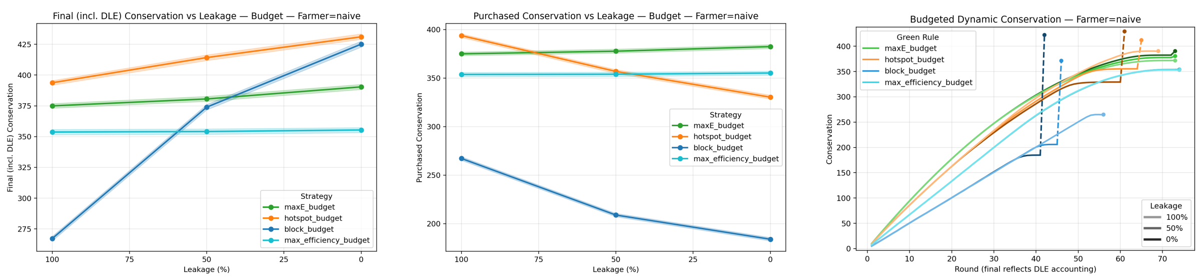

(a) Final (incl. DLE) Conservation **(b)** Purchased Conservation **(c)** Dynamic Trajectories

Figure 7: Budget World Green Strategies and Naïve Farmers: Conservation Scores

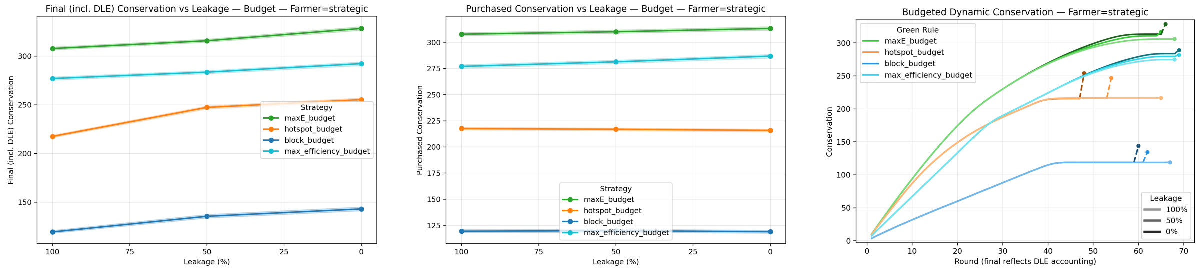

(a) Final (incl. DLE) Conservation **(b)** Purchased Conservation **(c)** Dynamic Trajectories

Figure 8: Budget World Green Strategies and Strategic Farmers: Conservation Scores

The Budget World decomposition helps to sharpen the distinction between purchased conservation (PC) and displacement/leakage effects (DLE). As in Claims World, DLE captures the indirect gains that occur when Farmer budget burn prevents BAU purchases. This decomposition makes clear that strategies can appear successful only because of residual crediting: for example, Hot Spot appears competitive against Naïve farmers in final conservation, but its advantage disappears in purchased outcomes.

Most importantly, Budget World allows us to introduce the *Max Efficiency* strategy, denoted by the turquoise lines above in Figures 7–8. One of the most striking and counter-intuitive findings of the paper is that Max Environment *always outperforms* Max Efficiency on purchased outcomes. Across 500 replications at each leakage level, Max Environment beats Max Efficiency in essentially 100% of cases, with a mean gap of 20–30 units of e. Moreover, in our simulations the PC gap is approximately invariant to leakage: the two strategies' PC curves are nearly parallel across all L, indicating that the advantage is driven by the statewise selection effect at disagreements rather than by residual crediting. This overturns the textbook knapsack heuristic that prioritizes e/a, and highlights a more general lesson: when cost is also a signal of rival demand, contesting expensive items is strictly better than avoiding them.

We formalize the core selection effect below. Theorem 4.1 shows that, under mild conditions, the value-greedy rule weakly dominates any alternative Green rule on expected Purchased Conservation (PC). In our simulations the resulting PC gaps are also nearly constant across L (the curves are close to parallel), consistent with the statewise selection mechanism rather than residual crediting; full proof is in Appendix A8.9.

Theorem 4.1 (PC dominance under nonnegative association (abbreviated; pointwise in leakage)). *Assume budget parity* ($s = 0.5$). *Let* (e_i, a_i) *be i.i.d. with continuous marginals and* nonnegative association *(e.g., TP_2/MTP_2; independence $e \perp a$ is the zero-association case). Suppose the Farmer's removal hazard is weakly increasing in an observable signal that is strictly increasing in a (covering both naïve rank-by-a and Strategic high-a targeting).*

For each fixed leakage level $L \in [0, 1]$, the value-greedy rule G^{env} that, in each round, purchases the affordable plot with the highest e weakly dominates *any alternative Green rule G on expected Purchased Conservation (PC), with* strict *inequality whenever G defers the current max–e affordable plot with positive probability. Because PC excludes the residual DLE crediting term, leakage does not enter the* statewise *selection argument: Lemma A8.8 holds for any L conditional on the state. The only channel through which leakage can influence PC is indirectly, through the evolution of budgets and hence the distribution of future opportunity sets; under "thick budgets" this indirect channel is asymptotically negligible, yielding global dominance for each fixed L (Theorem A8.9).*

Proof sketch. Couple G^{env} and G on the same instance and Farmer path. At the first disagreement G^{env} takes $k = \arg\max e$ while G takes j with $e_j < e_k$. With hazards increasing in an a-signal and nonnegative E–A association, k has (weakly) higher removal risk than j, so either k is lost before G returns (a PC loss) or G eventually buys k, leaving a gap of $e_k - e_j > 0$. Budget feedback terms vanish under thick budgets; under $\max_i a_i/B_g \to 0$, the budget–feedback remainder is asymptotically negligible, so the cumulative one–step gains yield expected-PC dominance in the limit. Since PC excludes the residual DLE crediting term, L does not enter the conditional one–step gain; any dependence of PC on L operates only through leakage-induced shifts in the state path, which vanish asymptotically under thick budgets. See Theorem A8.9 for the coupling and Proposition A8.11 for comparative statics. □

In Appendix A8 we provide a formal statement and proof sketch of this leakage-invariant dominance that, with a constructive counterexample and tiny grid equilibrium exercises, place the Budget World result on firm theoretical ground: the simulations are not simply numerical artifacts, but the manifestation of a deeper game-theoretic mechanism. Furthermore, although our conservation application emphasises that ecological and agricultural values may be weakly correlated in practice, the Budget World dominance result does not rely on this feature. Theorem 4.1 only assumes i.i.d. (e_i, a_i) with continuous marginals and non-negative association, and our robustness exercises with positive/negative correlation preserve the MaxE > Max-Eff hierarchy. Taken together, these results demonstrate that the Claims World insights are not artifacts of stylization, but robust features of adversarial conservation contests, and that efficiency-weighted heuristics do not necessarily outperform simpler value-maximizing rules in dynamic, strategic environments.

4.1.1 Budget World: Welfare Loss outcomes

Figures 9–10 confirm the *welfare paradox* observed in *Claims World*, persists *Budget World*: we observe that for some Green strategies welfare losses rise as leakage falls.

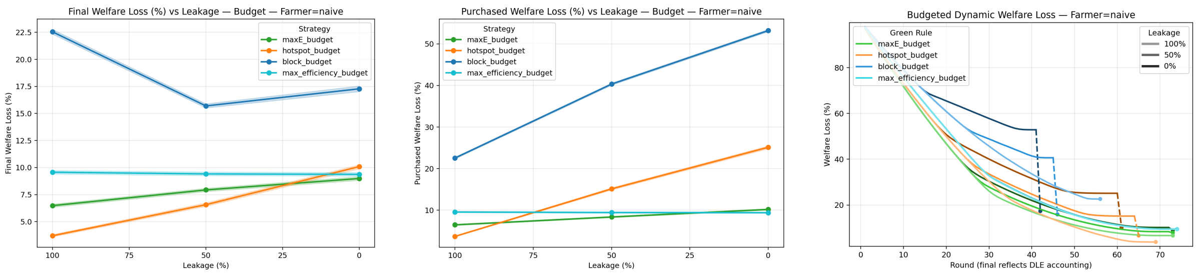

(a) Final (incl. DLE) Welfare Loss (%) **(b)** Purchased Welfare Loss (%) **(c)** Dynamic Trajectories

Figure 9: Budget World Green Strategies and Naïve Farmers: Social Welfare Loss (%)

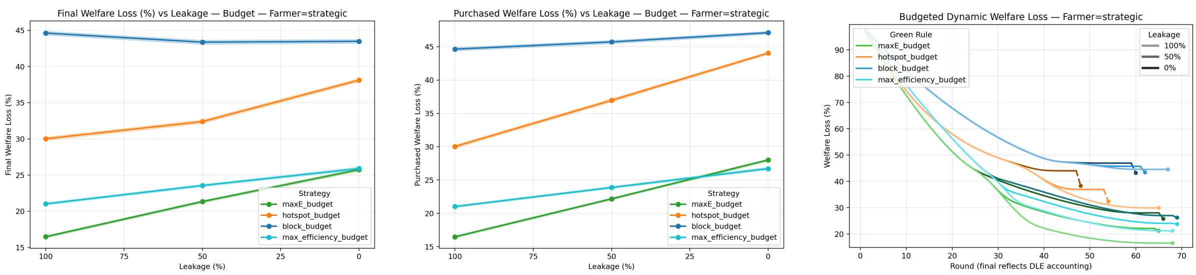

(a) Final (incl. DLE) Welfare Loss (%) **(b)** Purchased Welfare Loss (%) **(c)** Dynamic Trajectories

Figure 10: Budget World Green Strategies and Strategic Farmers: Social Welfare Loss (%)

5 Empirical Case Study: Conservation Strategies in Bolivia

Bolivia provides an insightful case study of the dynamic and contested nature of conservation decisions. With an extraordinary level of biodiversity and significant pressure for economic development, the history of Bolivian conservation efforts illustrate the trade-offs and strategic interactions that our dynamic adversarial framework addresses. Starting with the first national park in 1939, the land area under Protected Areas (PAs) status in Bolivia has gradually increased to 35.4 million hectares today, 32% of the national territory. Appendix section A9 provides short history of protected areas in Bolivia and describes the construction of our high-resolution country-wide dataset of the *potential* annual dollar values per hectare (USD/ha/year) from both conservation and agriculture (Andersen et al., 2023, 2024). First, in section 5.1 we construct a time series (in 5-year increments) of the relative environmental and agricultural quality and extent of land newly converted to both protected and developed areas, allowing us to examine historical patterns in land allocation through the lens of our theoretical framework. Second, in section 5.2 we depart from stylized grids and partition Bolivia into 390 hydrologically coherent Planning Units (PUs), calculating an environmental and agricultural score for each PU. This allows us to explore how our Green heuristic strategies perform when development and conservation values reflect patterns from real bio-physical data.

5.1 Conservation and agricultural expansion in Bolivia

The computational framework presented in this paper points to empirically novel summary statistics that may be informative. In particular, by comparing the average agricultural and environmental values of developed and protected land, we reveal shifts in Bolivian conservation practice over time, and by plotting time trends we can relate these shifts to periods of increased and decreased development pressure. Appendix section A9 provides short history of protected areas in Bolivia and describes the construction of the dataset. Our high resolution time series land use data combined with the potential conservation and agricultural values described in section A9.2 allows us to examine patterns in land allocation through the lens of our theoretical framework. Specifically, Figure 12 consolidates data on both the growth of agricultural and protected areas and their relative quality in terms of potential output and conservation, respectively, in 5-year increments. The area of new land allocated to Protected Areas is indicated by the height of upward bars, while new land allocated to Agriculture is shown by the downward bars. Each bar has both green and brown shades to indicate the average conservation and agricultural potential, respectively, of the newly allocated land - darker shades represent higher average values and lighter shades represent lower average values. For example, figure 12 shows that in 1986–1990 almost 4 million ha was allocated to Protected Area status, and just over 1 million ha was converted to agricultural land.

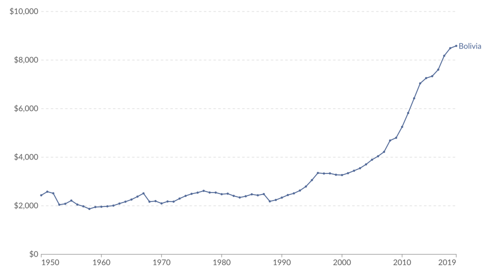

Figure 11: Bolivian GDP per capita, 1950–2019
(constant 2017 USD adjusted for cost of living)

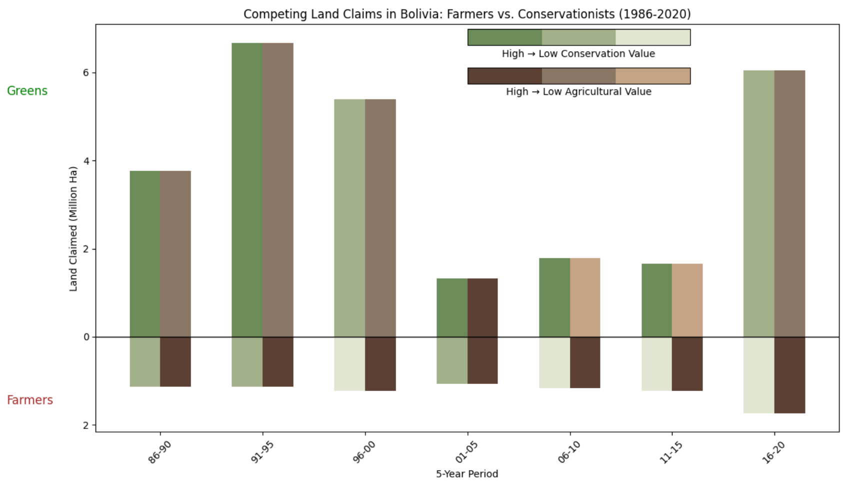

Figure 12: Agricultural & Conservation Potentials and Area of New Land allocated to Agriculture and Conservation, 5-year increments

Comparing the potential agricultural values (shades of brown) in Figure 12 we can see that the land allocated to agriculture is a darker brown (higher potential) than that allocated to conservation. Comparing the potential conservation values (shades of green) we see that the average conservation value of land allocated for Protected Area status is much higher (darker) than that allocated to agriculture. For example, the period 1986–1990 shows patterns where newly protected areas had high conservation values while agricultural expansion targeted high agricultural-value land with lower conservation value—a pattern consistent with each side pursuing their respective values. Throughout the sample period, agricultural expansion consistently targets high agricultural-potential land that tends to have lower conservation value—a pattern consistent with profit maximization. The patterns in conservation targeting appear more varied.

Figure 11 shows the evolution of GDP per capita over the same period. Comparing this with Figure 12 we can observe that the timing of protected area expansion correlates notably with economic cycles: major expansions coincide with economic downturns (late 1980s, 2019–2023) when extractive pressure decreases and international conservation funding may be relatively more influential.

The characteristics of newly protected land also show interesting temporal variation. Early protected areas (pre-1995) were predominantly in regions of very high conservation value. During the economic boom of 2001–2005, newly protected areas appeared more frequently near agricultural frontiers—a pattern consistent with threat-based targeting. This can be observed in Supplementary video A9.2 where new Protected Areas during this period often appear within areas of existing agricultural expansion. From 2006 onwards, the pattern shifts back toward high conservation value areas, though with more variation.

The documented failure of Protected Area status at Laguna Concepción to prevent agricultural encroachment (see section A9) illustrates how enforcement challenges may create a negative correlation between agricultural pressure and realized conservation outcomes. Appendix section A3 explores how such enforcement weakness might affect strategy performance in our simulations.

In sum, the patterns observed in Bolivia—high conservation value targeting during most periods, shifts toward threat-based targeting during economic booms, and challenges with enforcement at high-agricultural-value sites— illustrate a real-world context for considering our simulation results and underscore how viewing conservation through the lens of contested, sequential decisions can reveal patterns obscured by aggregate statistics. The correlation between conservation expansion and economic downturns, the apparent shifts in targeting approaches over time, and the vulnerability of high-agricultural-value sites like Laguna Concepción all become visible when agricultural and environmental values are examined jointly. The exercise illustrates how our framework provides a structured vocabulary to help interpret the complex dynamics of conservation practice.

5.2 Bridging the gap from theory to landscape: conservation strategies on a Bolivian 'board'

In section 5.1 we used our high resolution estimates of agricultural and environmental potential descriptively to characterise how historical protected areas and agricultural expansion have been sited over time.To test whether the analytical results from the computational simulations hold on a real landscape, we now use this data to transplant our Budget World contest onto the actual geography of Bolivia, moving the analysis from a stylized grid to a board characterized by realistic heterogeneity that reflects patterns from real bio-physical data.

5.2.1 Building a Bolivian planning board

We begin by aggregating the Bolivian land value pixels into a tractable but ecologically meaningful set of planning units (PUs). We use the nested HydroBASINS units for South America and clip them to Bolivia, selecting level–7 basins and splitting a handful of very large units with level–8 subdivisions. This yields hydrologically coherent PUs whose sizes range from a few hundred to a few tens of thousands of square kilometres. Using hydrological basins serves two purposes: first, it respects ecological boundaries; watershed integrity is often the minimum viable unit for conservation. Second, it introduces a realistic distributional and correlational structure to the 'grid.' We then remove tiny sliver units, PUs with no agricultural potential (no valid pixels in the agricultural raster), and one additional PU with very low environmental and agricultural scores (to ensure an even number of PUs), leaving a final board of $J = 390$ playable basins. The resulting Bolivian PUs are shown in Figure 13(a), overlayed with the 2024 extent of actual Protected Areas.

For each PU we then compute the mean environmental value and mean agricultural value (the μ scores), together with a set of "tail" statistics that capture how much of the PU lies in the national top deciles of the pixel–level distributions (90th, 95th and 99th percentiles). These are combined into composite scores $S_j^{(e)}$ and $S_j^{(a)}$ that summarise, respectively, each basin's *conservation attractiveness* and *agricultural opportunity cost* on a unit–free $[0, 1]$ scale (formal definitions are provided in Appendix A9.3):

- **Environmental Value ($S_j^{(e)}$):** We utilize the pixel-level ecosystem service valuation from Andersen et al. (2024). To aggregate this to the basin level, we compute a composite score that weights the mean value ($\mu_j^{(e)}$) and a set of "Tail-Averaged Percentile Scores" (TAPS) that capture the density of high-value pixels (90th, 95th, and 99th national percentiles). This ensures that basins containing small but intense "hotspots" of biodiversity are ranked highly, even if their average value is moderate.
- **Agricultural Price ($S_j^{(a)}$):** Similarly, we derive an agricultural score based on the potential net revenue maps from Andersen et al. (2023).

Both scores are normalized to a $[0, 1]$ scale based on their percentile ranks among all PUs, standardizing the "cost" of a PU as a measure of its opportunity cost relative to the national average. Thus we move the analysis from a stylized grid to a board characterized by realistic heterogeneity: environmental values are heavy-tailed, agricultural values follow complex soil and transport gradients, and the two values are positively correlated (with $\rho = .29$). The distribution of agricultural and environmental values is displayed in Figures 13(b) and 13(c).

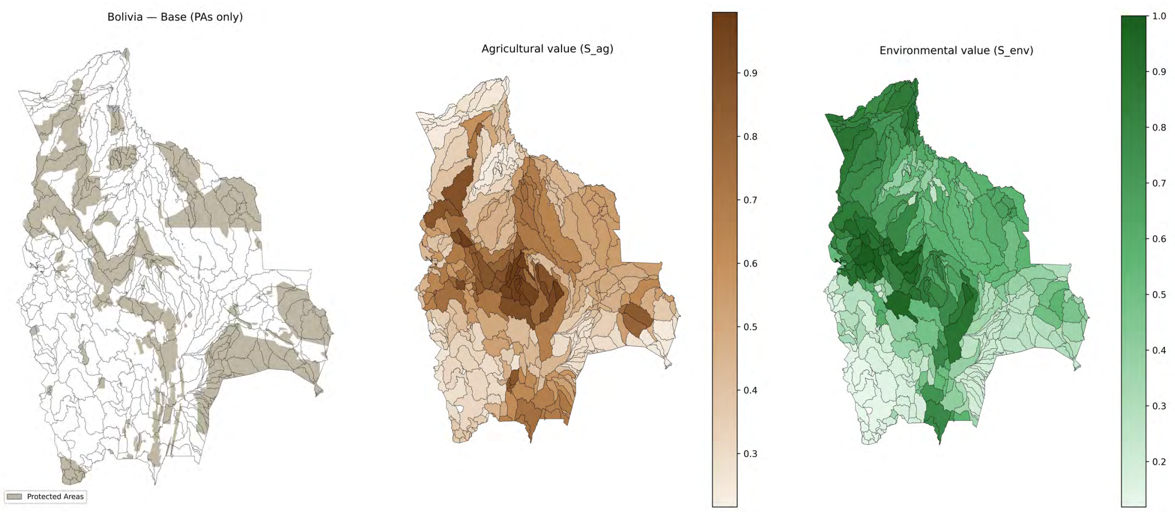

(a) PUs with PAs **(b)** Agricultural Values Heatmap **(c)** Environmental Values Heatmap

Figure 13: Bolivian PU board with 2024 Protected Areas and Heatmaps

We then run the Budget World contest on the fixed Bolivia board, simulating the competition between our Green strategies against both Naïve and Strategic developers across the full range of leakage parameters. As the Bolivia board is a single set of values, we obtain deterministic results.

The Bolivia board results for Claims World contests in Figures 14 - 17 closely mirror the main Monte Carlo patterns, despite being generated on a single, fixed landscape with realistic value distributions rather than on i.i.d. synthetic grids.

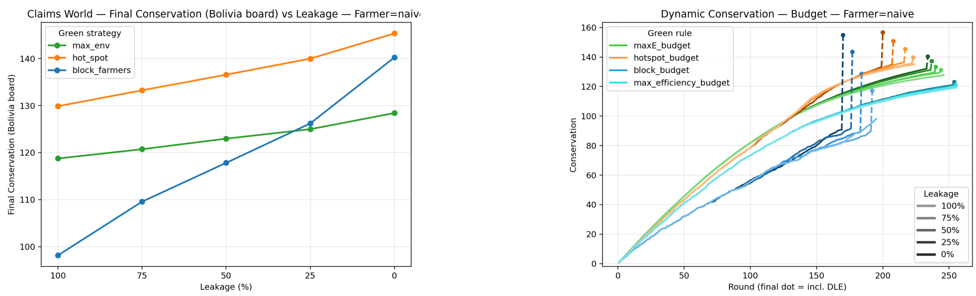

(a) Final Outcomes (incl. DLE) **(b)** Dynamic Trajectories

Figure 14: Claims World on a Bolivian Board with Naïve Farmers: Conservation Scores

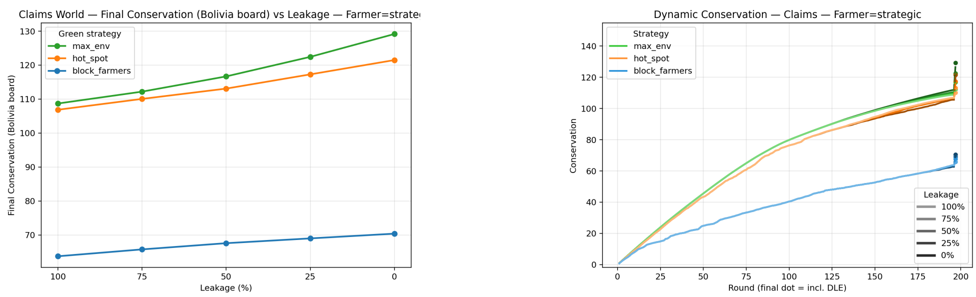

(a) Final Outcomes (incl. DLE) **(b)** Dynamic Trajectories

Figure 15: Claims World on a Bolivian Board with Strategic Farmers: Conservation Scores

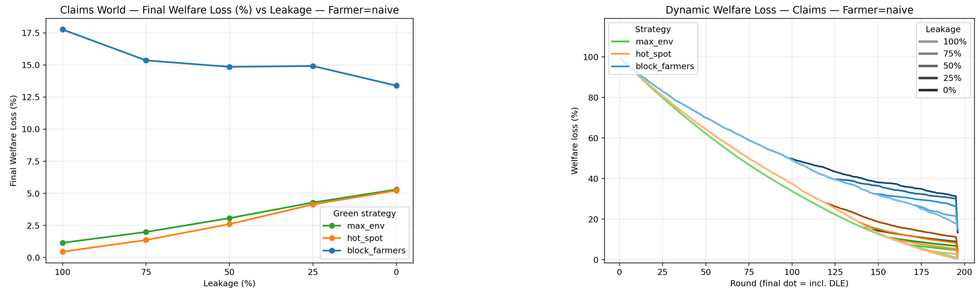

(a) Final Outcomes (incl. DLE) **(b)** Dynamic Trajectories

Figure 16: Claims World on a Bolivian Board with Naïve Farmers: Welfare Loss (%)

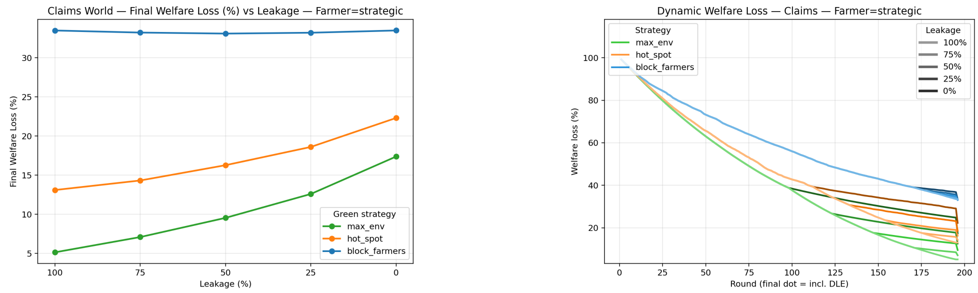

(a) Final Outcomes (incl. DLE) **(b)** Dynamic Trajectories

Figure 17: Claims World on a Bolivian Board with Strategic Farmers: Welfare Loss (%)

Claims World conservation outcomes on the Bolivia board show the same qualitative ranking of Green strategies we observe from the more stylized Monte Carlos (Figures 1 - 2 and Figures 5 - 6). Against Naïve Farmers (Figure 14), Hot Spot attains the highest final conservation score at all leakage levels, with Max Environment second and Block Farmers consistently worst. The dynamic trajectories in Figure 15(b) reveal the now-familiar mechanism: Max Environment initially produces the steepest Pure Strategy Effect (PSE), as it always takes the highest environmental-value basin, but Hot Spot gradually catches up and overtakes it as leakage falls, because it repeatedly secures high-e, high-a basins that would otherwise be developed. The end-of-game jumps again reflect Displacement–Leakage Effects (DLE) from residual unclaimed basins; as on the Monte Carlo grids, Hot Spot's final advantage over Max Environment comes primarily through larger DLE when leakage is low, not because it targets uniformly better sites throughout the game.

With Strategic Farmers (Figure 16) the ranking reverses exactly as in the synthetic grids:

Max Environment now strictly dominates Hot Spot in final conservation at all leakage levels, while Block Farmers collapses. The dynamic trajectories in Figure 16(b) show that once Farmers actively seek high-environmental basins, any deviation from a pure e-first rule allows them to capture contested PUs that Hot Spot or Block Farmers neglect. This echoes the main text result that value-first play becomes essential when developers behave strategically: the BAU sets of Strategic Farmers are more heavily loaded with high-e PUs, so Max Environment benefits both in PSE and in the DLE that accrues when those BAU plots are successfully blocked.

Figures 16 and 17 report welfare losses on the Bolivia board mirror the patterns from the Monte Carlos (Figures 5 and 6). Against Naïve Farmers, Max Environment and Hot Spot incur the lowest welfare losses across leakage levels, while Block Farmers produces much larger losses. A striking feature, which parallels the Monte Carlo results, is that welfare loss for Max Environment and Hot Spot *increases* as leakage falls: the opportunity set gradually tilts toward high-a sites, so successful blocking implies flipping more plots that optimally belong to agriculture. With Strategic Farmers, welfare losses are higher in level, but the ranking is unchanged: Max Environment is always best, Hot Spot intermediate, and Block Farmers worst. The welfare paradox persists on the real landscape: stronger leakage control can raise misallocation costs by forcing high-agricultural-value basins into conservation without commensurate ecological gains.

Budget World results on the Bolivia board for conservation and welfare (Figures 18 - 19) likewise confirm the robustness of the main simulation findings (the remaining results on additionality are similarly consistent with the Monte Carlos but not presented here as the patterns are identical to those for Conservation).

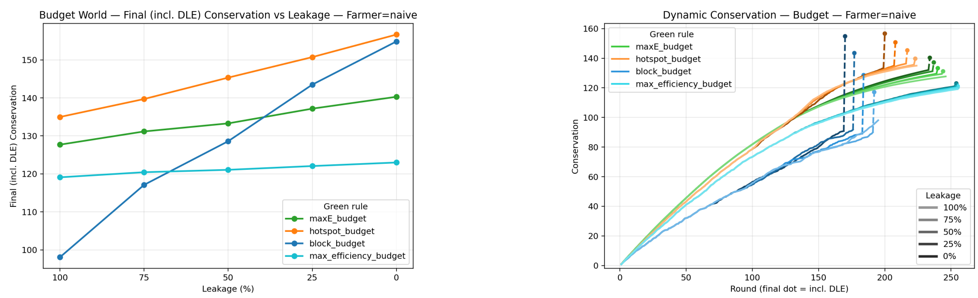

(a) Final Outcomes (incl. DLE) **(b)** Dynamic Trajectories

Figure 18: Budget World on a Bolivian Board with Naïve Farmers: Conservation Scores

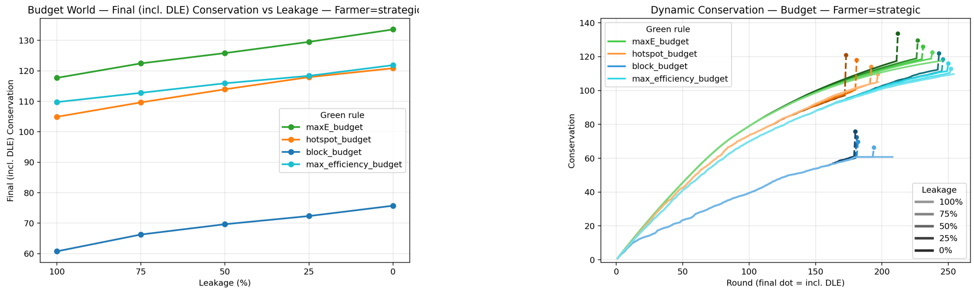

(a) Final (incl. DLE) Conservation **(b)** Dynamic Trajectories

Figure 19: Budget World on a Bolivian Board with Strategic Farmers: Conservation Scores

In Budget World, against Naïve Farmers, Hotspot_budget achieves the highest final conservation at most leakage levels, with MaxE_budget close behind; both substantially outperform Block_budget and MaxEfficiency_budget. The dynamic trajectories again show that Hotspot_budget's apparent edge comes largely from DLE when leakage is low: its Purchased Conservation (PC) path is only slightly above that of MaxE_budget, and the gap in final outcomes is driven by the residual crediting of undeveloped basins. Against Strategic Farmers, the pattern aligns even more tightly with the stylized Budget World Monte Carlo results: MaxE_budget clearly dominates on conservation, MaxEfficiency_budget is uniformly inferior, and Block_budget performs poorly. The PC curves for MaxE_budget and MaxEfficiency_budget are nearly parallel across leakage, with a roughly constant gap, reproducing on the Bolivia board the leakage-invariant selection effect highlighted in Theorem 4.1.

Figures 20 and 21 show the spatial distribution of the chosen PUs by MaxE_budget and Hot Spot_budget against both Naïve and Strategic Farmers.

(a) against Naïve Farmers **(b)** against Strategic Farmers

Figure 20: Budget World Max Environment Strategy Conserved PUs (leakage=100%)

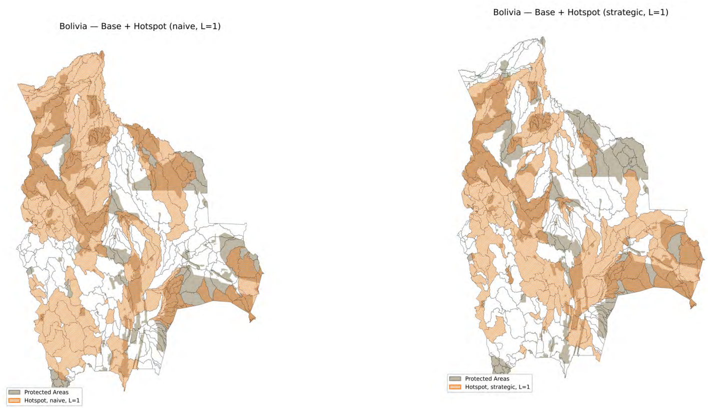

(a) against Naïve Farmers **(b)** against Strategic Farmers

Figure 21: Budget World Hot Spot Strategy Conserved PUs (leakage=100%)

Overall, then, the Bolivia board behaves as a realistic instantiation of the stylized games. Despite heavy-tailed environmental scores, spatial clustering of high-value basins, and positive but imperfect correlation between $S^{(e)}$ and $S^{(a)}$, the core comparative statics are unchanged: Block Farmers performs badly; Hot Spot can edge ahead against naïve developers but collapses against strategic ones; and the value-greedy Max Environment strategy remains the most reliable performer on both conservation and welfare, particularly when developers are strategic. In Budget World, the simple MaxE_budget rule continues to dominate the canonical ratio-greedy MaxEfficiency_budget rule even on this real landscape.

5.2.2 Green Strategies and Marxan planning

Many practical conservation planning exercises today are built around static optimization tools, most notably Marxan (Ball et al., 2009; Watts et al., 2009) which solve static target-based reserve design problems under a non-reactive (passive-developer) assumption. These systems formalize the planner's problem as a "minimum set problem": selecting a subset of planning units that meets ecological representation targets at the lowest possible cost (Ando et al., 1998b). Acquisition costs are treated as frictional constraints to be minimized, and threats are generally treated as exogenous probabilities. The resulting "optimal" portfolio assumes that selecting a site for protection removes it from the market without triggering a strategic response from developers. This logic has been adopted globally, with ratio-based cost-effectiveness heuristics used to design reserve networks in over 100 countries as of 2009 (Watts et al., 2009).

In applied settings, a common simplified analogue is a static cost-effectiveness ranking that prioritizes planning units by benefit–cost ratios—essentially a static "ratio-greedy" strategy. However, our computational adversarial simulations found that the "Ratio-Greedy" heuristic (Max Efficiency) is systematically dominated by "Value-Greedy" rules when costs act as signals of rival demand. To test whether this "Knapsack Reversal" holds on a real landscape, we first calculate the outcome of a widely used static cost-effectiveness heuristic. To align the static benchmark with our outcome metric (total ecological value purchased under a fixed budget), we therefore use a knapsack-form static planner problem (rather than Marxan's canonical minimum-set target formulation); the common element is the static, non-reactive assumption. Specifically, we provide the Green agent with a budget B_G and allow them to purchase PUs unopposed, solving:

$$\max_{x_j \in \{0,1\}} \sum_j S_j^{(e)} x_j \quad \text{s.t.} \quad \sum_j S_j^{(a)} x_j \leq B_G.$$

We implement this via the greedy algorithm, selecting PUs in descending order of their benefit-cost ratio $(S_j^{(e)}/S_j^{(a)})$. We refer to this as a *Marxan-style* static benchmark because the heuristic mirrors the non-reactive, budget-constrained site-selection logic common in systematic conservation planning. As a robustness check, we also solve the passive-developer static portfolio exactly as a 0–1 integer linear program (ILP) and get nearly identical results: both select 209 units, with an overlap of 208 units (Jaccard similarity = 0.9905). The ILP solution differs by a single swap and increases total purchased ecological value by $\approx 0.02\%$. Hence all qualitative conclusions are unchanged.

We then run the Budget World contest on the fixed Bolivia board, simulating the competition between our Green strategies against both Naïve and Strategic developers across the full range of leakage parameters. The results, presented in Figures 22 and mapped out in figure 23, yield two striking and related insights regarding the performance of Marxan conservation strategies in a real-world contested environment.

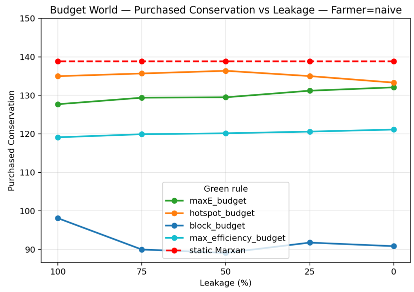

(a) Purchased Conservation and Naïve Farmers

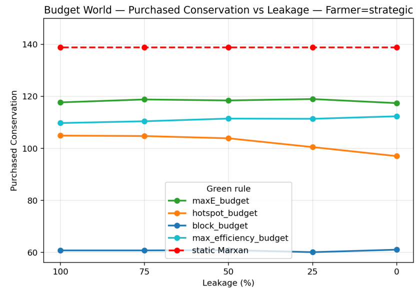

(b) Purchased Conservation and Strategic Farmers

Figure 22: Budget World on a Bolivian Board: Final Purchased Conservation

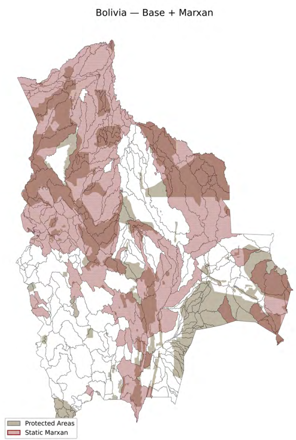

(a) Static Marxan-style

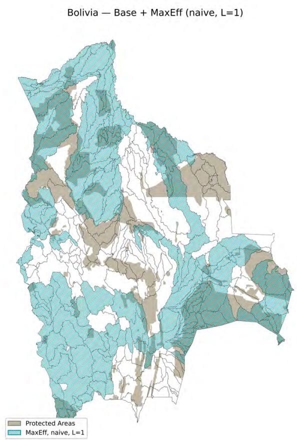

(b) Dynamic Max Efficiency Naïve

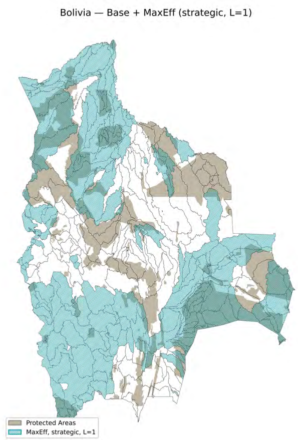

(c) Dynamic Max Efficiency Strategic

Figure 23: Static Marxan-style Benchmark and Dynamic Max Efficiency Conserved PUs (Budget World, leakage=100%)

First, we observe a significant gap between the Static Marxan-style Benchmark (red dashed line) and the realized outcomes of the Max Efficiency strategy (cyan line) under dynamic adversarial play. Across all leakage levels, the static plan overestimates realized conservation value by approximately 15–20%. This quantifies what we term the "disappointment gap"—a dynamic dimension of the broader 'implementation gap' (Knight et al., 2008)—where a planner using standard static tools to optimize a portfolio will systematically underperform expectations because the efficiency heuristic targets low-cost sites that are rarely contested, while failing to secure the high-value sites that developers actively target.

Second, the simulations confirm that the "Knapsack Reversal" identified in our stylized models is robust to real-world geography. The simple **Max Environment** strategy (green line) consistently outperforms the **Max Efficiency** strategy (cyan line) in Purchased Conservation (PC). Even on a jagged, heterogeneous map where costs and values are correlated, the strategy of contesting high-value sites ("Value-Greedy") secures a superior portfolio to the strategy of seeking bargains ("Ratio-Greedy").

6 Discussion

What are the most effective conservation strategies for achieving environmental objectives in adversarial settings? Most studies evaluate strategies from a planner's perspective, treating threats as exogenous. The limited non-cooperative work to date (e.g. Angelsen, 2001; Colyvan et al., 2011) highlights that conflict leads to suboptimal outcomes but offers little guidance on which siting rules perform best. Our contribution fills this gap by comparing widely-practised conservation strategies in a computational framework where developers actively respond to conservation pressure.

A central innovation of our framework is to decompose the performance of each strategy into two analytically distinct components: the Pure Strategy Effect (PSE), reflecting the intrinsic quality of the plots chosen by a strategy, and the Displacement–Leakage Effect (DLE), reflecting the gains from displacing developers' preferred sites when leakage is incomplete. This decomposition shows that strategies that appear effective in aggregate may achieve their gains in very different ways, with important implications for how robust those gains are under different institutional conditions. We explore these effects in two 'Worlds': in *Claims World* we abstract from cost to isolate the implications of rivalry, and in *Budget World* we introduce a continuous procurement mechanism under budget constraints.

Our simulations yield several striking and robust findings. In Claims World, the direct link between threat-weighting and additionality—so central in planner-world models—breaks down once developer behaviour is endogenised. In our decomposition, the PSE of threat-targeting is often weak, and when leakage is high and developers can re-target freely, the DLE evaporates, leaving little more than random conservation. Strategies such as Block Farmers therefore perform poorly, even though they appear to have high additionality when threats are treated as fixed.

The PSE/DLE decomposition also sheds light on a surprising paradox: reducing leakage does not uniformly improve outcomes. On the contrary, it can increase welfare losses, as high-development-value plots are shifted into conservation without commensurate ecological gain. At lower levels of leakage, the PSE of Hot Spot and Max Environment tilts the opportunity set toward contested high-a plots, increasing the chance of misallocation in later rounds. Reduced leakage may also boost DLE, but by forcing agriculturally valuable residual plots into conservation this mechanism can simultaneously raise welfare costs. These patterns appear robust across both Claims and Budget Worlds, and in our Bolivia case study we show that they translate closely onto the organically constructed Bolivia board, where agricultural and environmental values are derived from real bio-physical data.

Crucially, the relative ranking of strategies depends on developer behaviour. Against naïve profit maximisers, Hot Spot targeting can secure higher final conservation scores, but against strategic developers Max Environment dominates in both conservation and welfare terms. The decomposition again clarifies why: Hot Spot's advantage comes primarily from DLE—it extracts value by blocking developers in a few high-risk frontiers—whereas Max Environment derives its strength from consistently high PSE, selecting intrinsically valuable sites across the board. The attractiveness of strategies that focus squarely on ecological value thus emerges as a central lesson of the analysis, particularly once we allow for adaptive developers.

Budget World analysis confirms the robustness of the patterns from Claims World and reveals our most striking result: a dynamic "Knapsack Reversal." We find that the standard cost-effectiveness heuristic—maximise environmental value per dollar (e/a)—is systematically dominated by the simple value-greedy rule (e). Intuitively, the ratio-greedy rule economises by buying cheap "bargains" that are rarely contested. However, in this two-space setting where prices signal rival demand, avoiding expensive plots implies retreating from the contest. Value-greedy rules succeed because they contest the high-a region where developers concentrate. We formalize this in Theorem 4.1: under non-negative e–a association, value-first procurement weakly dominates cost-effectiveness heuristics on purchased conservation. This suggests that static Marxan-style portfolios may systematically disappoint in adversarial settings because they optimize for a landscape that does not react.

Crucially, the 'Knapsack Reversal' and the dominance of value-greedy rules are not artifacts of the computational framework, but reflect fundamental properties of the adversarial game structure confirmed by our formal analysis. More generally, our simulation patterns are consistent with best-response logic in tractable limiting cases (Appendix A8). When developers can freely re-target ($L \to 1$), blocking confers no dynamic benefit and the Green best reply reduces to a myopic *preempt-or-value* rule, under which pure threat-chasing collapses toward random performance. Heuristically, this equilibrium strategy can be interpreted as "Maximise Environment, except when the plot about to be claimed by developers also has the highest environmental value among the remaining opportunities"—in other words, Maximise Environment with judicious use of Hot Spot. When re-targeting is limited ($L < 1$), the Green objective acquires a "business-as-usual" (BAU) bonus for plots that developers would otherwise take, yielding a balanced e-first threshold rule: focus on ecological value, but with a limited tilt toward genuinely contested high-e, high-a sites.

These insights are validated by our empirical case study. By transplanting the contest to a planning board derived from Bolivian biophysical data—characterized by heavy-tailed environmental distributions and spatially correlated values—we confirm that these patterns are not artifacts of stylized grids. The simulation reveals a quantifiable gap between the static Marxan benchmark and realized dynamic outcomes. Even on this complex real-world landscape, the Knapsack Reversal holds: dynamic value-greedy strategies outperform the ratio-greedy heuristic, and the robustness of Max Environment against strategic developers persists.

Conceptually, the paper thus adds a new piece to the theory of adversarial allocation: in two-space contests where assets carry one value for a planner and another for a rival, and prices track the rival's value, rules that defer expensive high-value items in favour of cheap "bargains" are dominated on purchased outcomes. Our Budget World knapsack reversal is therefore best viewed as a more general mechanism that applies whenever cost doubles as a signal of rival demand. Conservation provides a particularly transparent instance of this structure, because ecological and development values are naturally distinct, but the same logic applies in any setting where social and rival values live in separate spaces - for example, where social housing funds seek to acquire high-opportunity properties and the planner's social value of long-run affordability is largely decoupled from the market price, but prices are set by private developers' returns. In such settings, our analysis suggests that value-first procurement rules that contest expensive, high rival-value assets may outperform cost-effectiveness heuristics that chase cheap bargains, whenever budgets are deployed against an active rival.

The policy implications are clear. Conservation and green finance initiatives should be cautious about relying too narrowly on threat-based additionality metrics or cost-effectiveness mandates, which risk fragility in the face of market adaptation. Robust conservation requires strategies that anticipate adaptation, not just static threats. A balanced portfolio—investing not only in high-threat frontier lands but also in ecologically irreplaceable sites—generates more durable conservation gains and prevents the strategic erosion of environmental value.

References

Abbitt, R., Scott, J., and Wilcove, D. (2000). The geography of vulnerability: Incorporating species geography and human development patterns into conservation planning. *Biological Conservation*, 96:169–175.

Allan, J. R., Venter, O., Maxwell, S., Bertzky, B., Jones, K. R., Shi, Y., and Watson, J. E. M. (2019). Hotspots of human impact on threatened terrestrial vertebrates. *Nature Communications*, 10(1):1–9.

Almond, R., Grooten, M., Juffe Bignoli, D., and Petersen, T., editors (2022). *Living Planet Report 2022 – Building a nature-positive society.* WWF, Gland, Switzerland.

Amin, A., Choumert-Nkolo, J., Combes, J.-L., Combes Motel, P., Kéré, E., Ongono-Olinga, J.-G., and Schwartz, S. (2019). Neighborhood effects in the brazilian amazônia: protected areas and deforestation. *Journal of Environmental Economics and Management*, 93:272–288.

Andam, K. S., Ferraro, P. J., Pfaff, A., Sanchez-Azofeifa, G. A., and Robalino, J. A. (2008). Measuring the effectiveness of protected area networks in reducing deforestation. *Proceedings of the National Academy of Sciences*, 105(42):16089–16094.

Andersen, L., Argondoña, F., and Calderón, D. (2024). *Aproximación al valor de las contribuciones de las Áreas Protegidas y los Territorios Indígenas Originarios Campesinos de Bolivia al bienestar humano a nivel local y global.* WWF y SDSN Bolivia. Borrador todavía no publicado.

Andersen, L. E., Argandoña, F., Choque Sunagua, S., Calderón Acebey, D. L., Inkinen, V., and Malky, A. (2023). Map of agricultural potential in bolivia. SDSN Working Paper No. 5/2023.

Ando, A., Camm, J., Polasky, S., and Solow, A. (1998a). Species distributions, land values, and efficient conservation. *Science*, 279:2126–2128.

Ando, A., Camm, J., Polasky, S., and Solow, A. (1998b). Species distributions, land values, and efficient conservation. *Science*, 279(5359):2126–2128.

Angelsen, A. (2001). Playing games in the forest: State-local conflicts of land appropriation. *Land Economics*, 77(2):285–299.

Armsworth, P. R., Daily, G. C., Kareiva, P., and Sanchirico, J. N. (2006). Land market feedbacks can undermine biodiversity conservation. *Proceedings of the National Academy of Sciences*, 103(14):5403–5408.

Aspelund, K. M. and Russo, A. (2023). Additionality and asymmetric information in environmental markets: Evidence from conservation auctions. Working Paper.

Baldi, G., Texeira, M., Martin, O., Grau, H., and Jobbágy, E. (2017). Opportunities drive the global distribution of protected areas. *PeerJ*, 5:e2989.

Ball, I. R., Possingham, H. P., and Watts, M. (2009). Marxan and relatives: Software for spatial conservation prioritisation. In Moilanen, A., Wilson, K. A., and Possingham, H. P., editors, *Spatial Conservation Prioritisation: Quantitative Methods and Computational Tools*, pages 185–195. Oxford University Press, Oxford.

Balmford, A., Gaston, K., and Rodrigues, A. (2000). Integrating costs of conservation into international priority setting. *Conservation Biology*, 14:597–605.

Brams, S. J. and Kaplan, T. R. (2004). Dividing the indivisible: Procedures for allocating cabinet ministries to political parties in a parliamentary system. *Journal of Theoretical Politics*, 16(2):143–173.

Budish, E. and Cantillon, E. (2012). The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *American Economic Review*, 102(5):2237–2271.

Ceballos, G., Ehrlich, P., Barnosky, A., García, A., Pringle, R., and Palmer, T. (2015). Accelerated modern human-induced species losses: entering the sixth mass extinction. *Science Advances*, 1(5).

Colyvan, M., Justus, J., and Regan, H. M. (2011). The conservation game. *Biological Conservation*, 144(4):1246–1253.

Delacote, P., Le Velly, G., and Simonet, G. (2024). Distinguishing potential and effective additionality of forest conservation interventions. *Environment and Development Economics*, 117:234–243.

Delacote, P., Palmer, C., Bakkegaard, R., and Thorsen, B. (2014). Unveiling information on opportunity costs in redd: who obtains the surplus when policy objectives differ? *Resource and Energy Economics*, 36:508–527.

Dinerstein, E., Joshi, A. R., Hahn, N. R., et al. (2024). Conservation imperatives: Securing the last unprotected terrestrial sites harboring irreplaceable biodiversity. *Frontiers in Science*, 2:1349350.

Dinerstein, E., Olson, D., Joshi, A. R., et al. (2017). An ecoregion-based approach to protecting half the terrestrial realm. *BioScience*, 67(6):534–545.

Dinerstein, E., Vynne, C., Sala, E., et al. (2020). A Global Safety Net to reverse biodiversity loss and stabilize earth's climate. *Science Advances*, 6(36):eabb2824.

Duncanson, L., Liang, M., and Leitold, V. e. a. (2023). The effectiveness of global protected areas for climate change mitigation. *Nature Communications*, 14:2908.

Engel, S., Pagiola, S., and Wunder, S. (2008). Designing payments for environmental services in theory and practice: An overview of the issues. *Ecological Economics*, 65(4):663–674.

Faith, D. P. (1992). Conservation evaluation and phylogenetic diversity. *Biological Conservation*, 61(1):1–10.

Fearnside, P. M. (2017). Deforestation of the brazilian amazon. In Shugart, H. H., editor, *Oxford Research Encyclopedia of Environmental Science*. Oxford University Press.

Ferraro, P. and Simpson, D. (2002). The cost-effectiveness of conservation payments. *Land Economics*, 78(3).

Fischetti, M., Ljubić, I., Monaci, M., and Sinnl, M. (2019). Interdiction games and monotonicity, with application to knapsack problems. *INFORMS Journal on Computing*, 31(2):390–410.

Forest Trends' Ecosystem Marketplace (2023). Paying for quality: State of the voluntary carbon markets 2023. Ecosystem Marketplace Insights Report, November 28, 2023.

Game, E. T., Kareiva, P., and Possingham, H. P. (2008). Threats matter: incorporating the threat of extinction into conservation prioritization. *Ecology Letters*, 11(12):1235–1241.

Geldmann, J., Manica, A., Burgess, N. D., Coad, L., and Balmford, A. (2019). A global-level assessment of the effectiveness of protected areas at resisting anthropogenic pressures. *Proceedings of the National Academy of Sciences*, 116(46):23209–23215.

Hansen, A. J., Barnett, K., Green, M., Harris, G., Keane, R., Walker, R., Gross, J., Phillips, L., Joly, K., Maxwell, B., and Rotella, J. J. (2020). Global protected areas and imperiled species: a focus on biodiversity hotspots under threat. *Nature Ecology & Evolution*, 4:1501–1506.

IPBES (2019). Global assessment report on biodiversity and ecosystem services.

Joppa, L. N. and Pfaff, A. (2011). Global protected area impacts. *Proceedings of the Royal Society B: Biological Sciences*, 278(1712):1633–1638.

Knight, A. T., Cowling, R. M., Rouget, M., Balmford, A., Lombard, A. T., and Campbell, B. M. (2008). Knowing but not doing: selecting priority conservation areas and the research-implementation gap. *Conservation Biology*, 22(3):610–617.

Knowles-Nelson Stewardship Program (2024). Knowles-nelson gets its day in court. https://knowlesnelson.org/knowles-nelson-gets-its-day-in-court/. Accessed: 2025-09-12.

Kovenock, D. and Roberson, B. (2010). Conflicts with multiple battlefields. In Garfinkel, M. R. and Skaperdas, S., editors, *Oxford Handbook of the Economics of Peace and Conflict*, pages 503–531. Oxford University Press, Oxford.

Langpap, C. and Wu, J. (2017). Thresholds, perverse incentives, and preemptive conservation of endangered species. *Journal of the Association of Environmental and Resource Economists*, 4(S1):S227–S259.

Leverington, F., Costa, K., Pavese, H., Lisle, A., and Hockings, M. (2010). A global analysis of protected area management effectiveness. *Environmental Management*, 46:685–698.

Lu, Y., Wang, H., Zhang, Y., Liu, J., Qu, T., Zhao, X., Tian, H., Su, J., Luo, D., and Yang, Y. (2023). Combining spatial–temporal remote sensing and human footprint indices to identify biodiversity conservation hotspots. *Diversity*, 15(10):1064.

Lueck, D. and Michael, J. A. (2003). Preemptive habitat destruction under the endangered species act. *Journal of Law and Economics*, 46(1):27–60.

MapBiomas Bolivia Project (2024). Mapbiomas bolivia collection 2.0: Land use/land cover and transitions 1985 to 2024. Accessed in March 2025.

Margules, C., Nicholls, A., and Pressey, R. (1988). Selecting networks of reserves to maximize biological diversity. *Biological Conservation*, 43:63–76.

Naidoo, R., Balmford, A., Ferraro, P. J., Polasky, S., Ricketts, T. H., and Rouget, M. (2006). Integrating economic costs into conservation planning. *Trends in Ecology & Evolution*, 21(12):681–687.

Navia, R. (2022). El planeta ha perdido a la laguna concepción. *Revista Nomades.* https://revistanomadas.com/el-planeta-ha-perdido-a-la-laguna-concepcion.

Newburn, D., Berck, P., and Merenlender, A. (2006). Habitat and open space at risk of land-use conversion: Targeting strategies for land conservation. *American Journal of Agricultural Economics*, 88(1):28–42.

Our World in Data (2025). Gdp per capita in constant international-$ [dataset]. Data from Feenstra et al. (2015), Penn World Table (2021), with major processing by Our World in Data. Retrieved March 13, 2025.

Pfaff, A., Robalino, J., Herrera, D., and Sandoval, C. (2015). Protected areas' impacts on brazilian amazon deforestation: examining conservation – development interactions to inform planning. *PLOS ONE*, 10:e0129460.

Pferschy, U., Nicosia, G., Pacifici, A., and Schauer, J. (2021). On the stackelberg knapsack game. *European Journal of Operational Research*, 291(1):18–31.

SDSN Bolivia (2025). Sdsn bolivia homepage. Accessed March 13, 2025.

Searchinger, T., Wirsenius, S., Beringer, T., and Dumas, P. (2018). Assessing the efficiency of land use changes for climate mitigation. *Nature*, 564(7735):249–253.

Silva, D. S. and Nunes, S. (2025). Leakage effects from reforestation: Estimating the impact of agricultural displacement for carbon markets. *Land*, 14(5):963.

Sims, K. (2014). Do protected areas reduce forest fragmentation? a microlandscapes approach. *Environmental and Resource Economics*, 58:303–333.

Tilman, D., Clark, M., Williams, D., Kimmel, K., Polasky, S., and Packer, C. (2017). Future threats to biodiversity and pathways to their prevention. *Nature*, 546(7656):73–81.

Venter, O., Fuller, R. A., Segan, D. B., Carwardine, J., Brooks, T., Butchart, S. H. M., Di Marco, M., Iwamura, T., Joseph, L., O'Grady, D., Possingham, H. P., Rondinini, C., Smith, R. J., Venter, M., and Watson, J. E. M. (2014). Targeting global protected area expansion for imperiled biodiversity. *Nature*, 516:143–146.

Watts, M. E., Ball, I. R., and Possingham, H. P. (2009). Marxan and relatives: supporting spatially explicit conservation decision making. In Moilanen, A., Wilson, K. A., and Possingham, H. P., editors, *Spatial Conservation Prioritisation: Quantitative Methods and Computational Tools*, pages 185–195. Oxford University Press, Oxford.

Wunder, S. (2015). Revisiting the concept of payments for environmental services. *Ecological Economics*.

Wünscher, T. and Engel, S. (2012). International payments for biodiversity services: Review and evaluation of conservation targeting approaches. *Biological Conservation*, 152:222–230.

Appendix

This Appendix includes a range of robustness and heterogeneity analyses as well as some formal theoretical results and background material for our Bolivian vignette. Specifically A2 examines positive correlation between agricultural and environmental values; A3 simulates weak legal enforcement via negative correlation; and A4 explores unequal political power by allocating more claims to Farmers. Appendix A5 considers heavy–tailed environmental draws, creating rare but extremely high–value "superstar" plots. Appendix A6 introduces an explicit spatial structure (S1) with paired plots generating spillovers in e, and shows in Appendix A6.1 how this embeds in a more general graph-based spillover model (S2).

Appendix A7 formally sets up the Budget World. Appendix A8 presents our core theoretical results. Appendix A8.1 connects the simulation patterns to equilibrium logic in tractable polar cases: at $L = 1$ (full re-targeting) optimal play reduces to a myopic *preempt-or-value* rule; at $L = 0$ (no re-targeting) the Green objective acquires a BAU "bonus," yielding an e-first threshold policy. Building on this, in Appendix sections A8.2 - A8.2.4 we establish our most striking theoretical and simulation result: in Budget World the textbook ratio-greedy e/a rule is systematically dominated by the simple e-maximizing rule, across all leakage levels and against both Naïve and Strategic Farmers. This leakage-invariant dominance reverses standard knapsack logic and provides a general principle for adversarial procurement. Appendix A8.3 presents a series of Tiny Grid equilibrium exercises that validate the core results.

Appendix A9 provides supplemental material on the history of protected areas in Bolivia and data construction for our empirical vignette. Finally, to facilitate replication and teaching we provide code and links to two publicly available tools in Appendix A10, including (1) a Monte Carlo simulator with documentation for replicating results and exploring alternative strategies; and (2) an interactive browser-based game where users can play as either team against fixed strategies played by the computer.

Table A1.1: Conservation Scores Under Different Strategies and Leakage Levels: The *Pure Strategy* and *Displacement-Leakage* Effects

Farmer Strategy	Green Strategy	Leakage Level (%)	*Pure Strategy* Effect	*Displacement-Leakage* Effect	Final Conservation Score
naive	Max Environment	0	337	17	354
naive	Max Environment	50	333	6	339
naive	Max Environment	100	331	0	331
naive	Block Farmers	0	252	126	378
naive	Block Farmers	50	252	60	312
naive	Block Farmers	100	252	0	252
naive	Hot Spot	0	344	63	408
naive	Hot Spot	50	354	29	383
naive	Hot Spot	100	361	0	361
strategic	Max Environment	0	301	34	336
strategic	Max Environment	50	294	12	306
strategic	Max Environment	100	290	0	290
strategic	Block Farmers	0	142	7	148
strategic	Block Farmers	50	142	3	144
strategic	Block Farmers	100	142	0	142
strategic	Hot Spot	0	237	22	259
strategic	Hot Spot	50	241	11	252
strategic	Hot Spot	100	244	0	244

A1 Claims World Baseline Conservation Achieved

A2 Positive Correlation between Agricultural and Conservation Values

Our case study of Bolivia suggests that agricultural and conservation potentials in practice may be positively correlated. Thus, below we plot the dynamic trajectories of outcome variables for different combinations of strategies when the grid is initialized with a positive correlation of 0.30 between Conservation values and Agricultural values (500 replications and 50 rounds).

Overall, with a positive correlation between agricultural and environmental values, we observe that conservation scores tend to be lower across all strategies, but the additionality scores of the Hot Spot and Block Farmers strategies tend to be slightly higher. The dynamic trajectory patterns are broadly similar, with the Conservation values of the Hot Spot *Pure Strategy* effect slightly stronger than in the zero correlation case.

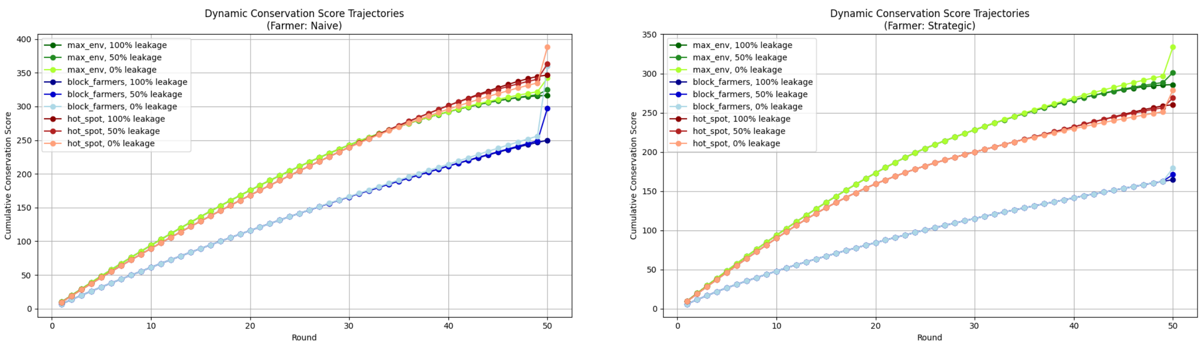

(a) Green Strategies and Naïve Farmers **(b)** Green Strategies and Strategic Farmers

Figure A2.1: Conservation Scores for $\rho = 0.3$

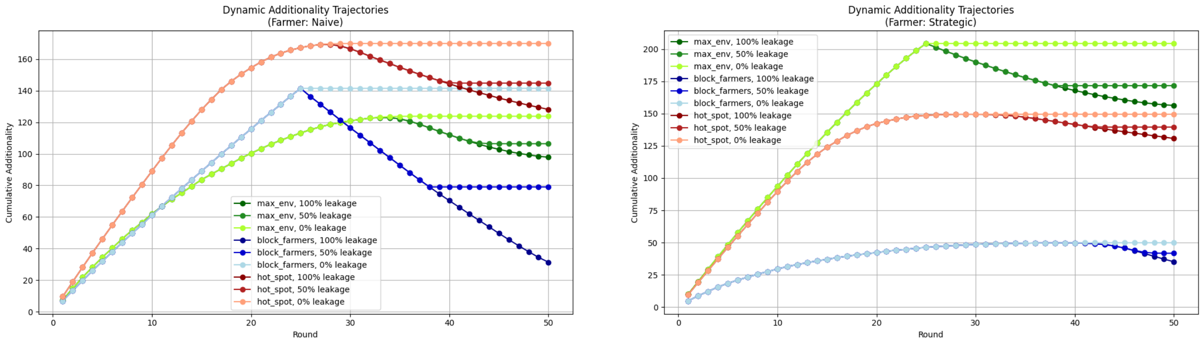

(a) Green Strategies and Naïve Farmers **(b)** Green Strategies and Strategic Farmers

Figure A2.2: Additionality for for $\rho = 0.3$

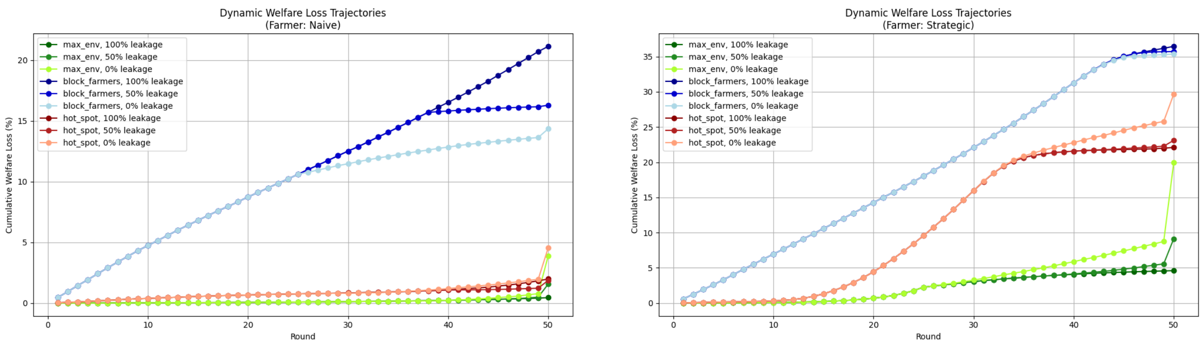

(a) Green Strategies and Naïve Farmers **(b)** Green Strategies and Strategic Farmers

Figure A2.3: Social Welfare Loss(%) for for $\rho = 0.3$

A3 Negative Correlation between Agricultural and Conservation Values

Our case study of Bolivia suggested that enforcement of protected status may be more challenging in areas under high threat of development, as illustrated by the experience of Laguna Concepción. In order to simulate this relationship we initialize the grid with a negative correlation between Agricultural and Environmental values - in those plots most desirable to the Farmers, realized Environmental values are lower due to increased difficulty of enforcement. Note that this approach explicitly differentiates between the effectiveness of leakage control (the *Displacement-Leakage* effect), which is still explored in the simulation by allowing for different levels of leakage, and a *Pure Strategy* effect when protection effectiveness is systematically reduced in Green claimed land of high agricultural value.

Thus, below we plot the dynamic trajectories of outcome variables for different combinations of strategies when the grid is initialized with a negative correlation of -0.30 between Conservation values and Agricultural values (500 replications and 50 rounds). Overall, with a negative correlation between agricultural and environmental values, the dynamic trajectory patterns are broadly similar, with the Conservation values of the Max Environment *Pure Strategy* effect slightly stronger than in the zero correlation case.

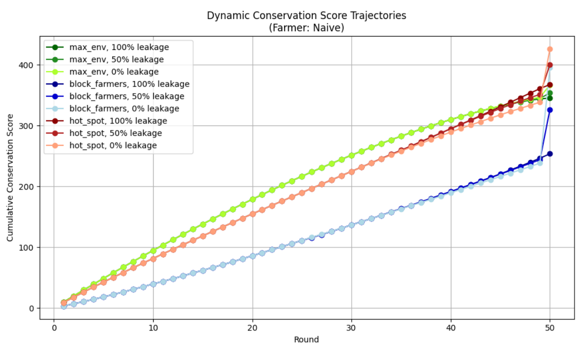

(a) Green Strategies and Naïve Farmers

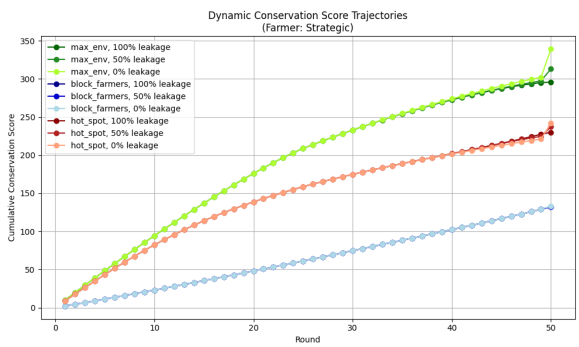

(b) Green Strategies and Strategic Farmers

Figure A3.1: Conservation Scores for ρ = -0.3

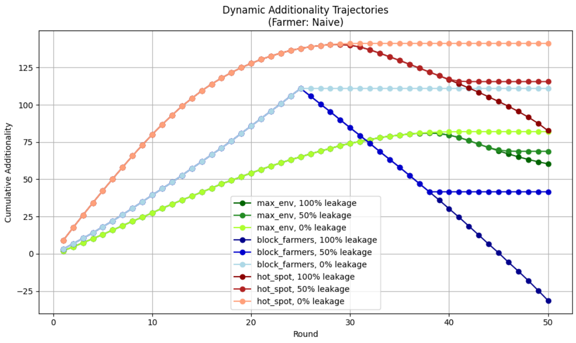

(a) Green Strategies and Naïve Farmers

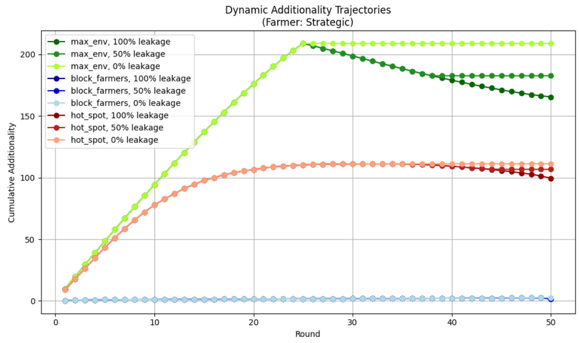

(b) Green Strategies and Strategic Farmers

Figure A3.2: Additionality for for ρ = -0.3

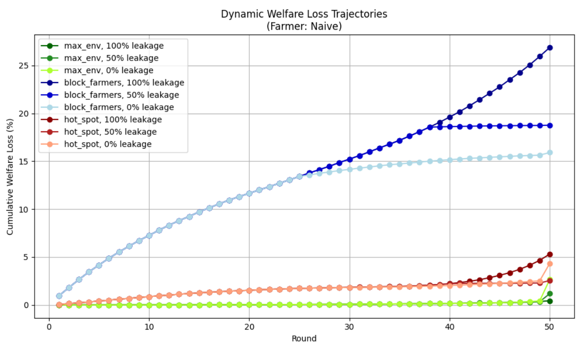

(a) Green Strategies and Naïve Farmers

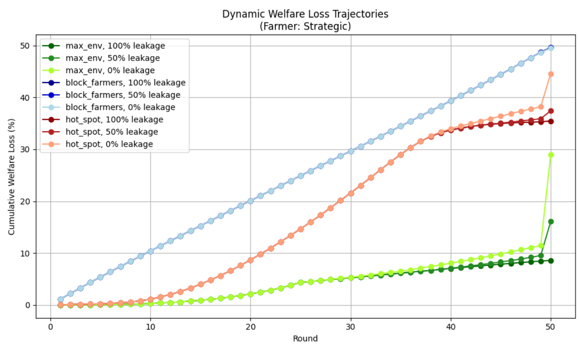

(b) Green Strategies and Strategic Farmers

Figure A3.3: Social Welfare Loss(%) for for ρ = -0.3

A4 Political Allocation in Favour of Farmers

Here we allocate 70% of the initial claims to the Farmers, exploring how the different Green strategies perform when conservation faces an economic or political disadvantage.

We present the dynamic trajectory paths below, and the final score is represented by the outcome in round 50. Since in the political allocation considered the Farmers will have excess claims and spend those to claim additional plots in the last round (while the Greens are allocated any remaining unclaimed plots), then in this scenario there are significant last-round adjustments to the final score across all three outcomes considered.

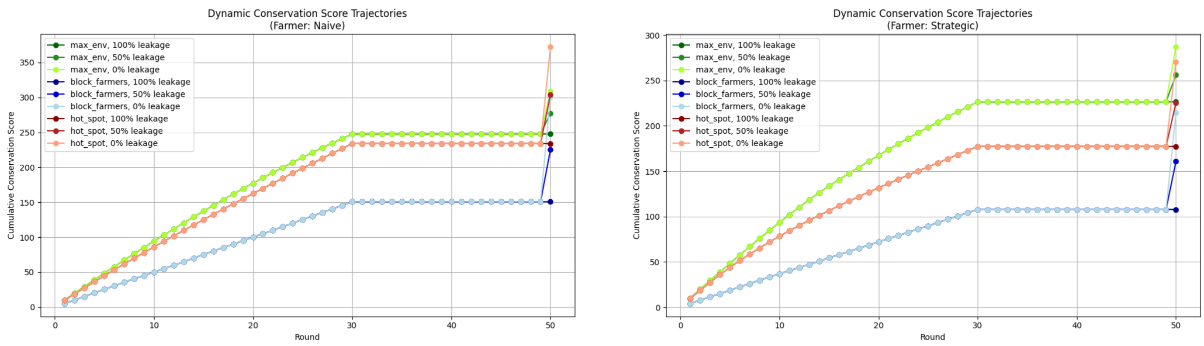

(a) Green Strategies and Naïve Farmers **(b)** Green Strategies and Strategic Farmers

Figure A4.1: Conservation Scores for Political Allocation to Farmers = 70%

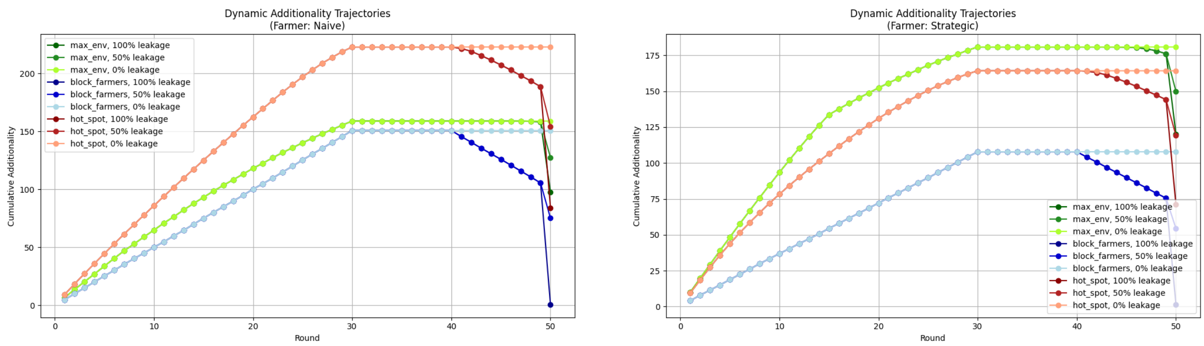

(a) Green Strategies and Naïve Farmers **(b)** Green Strategies and Strategic Farmers

Figure A4.2: Additionality for Political Allocation to Farmers = 70%

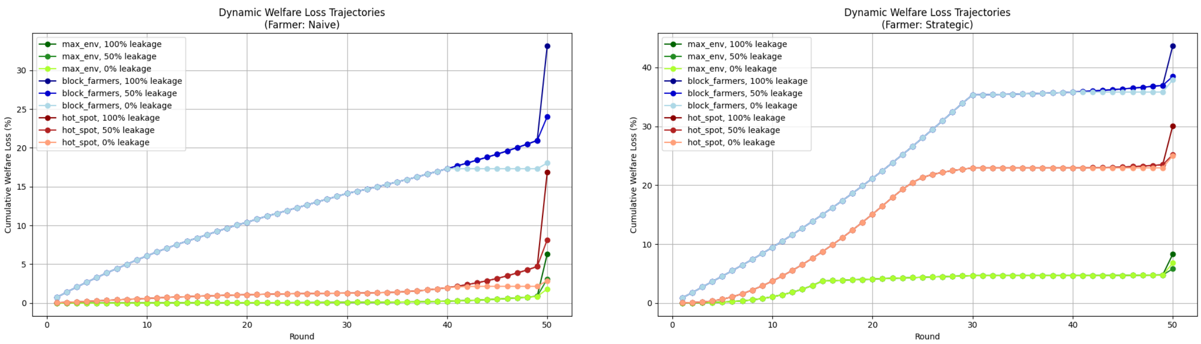

(a) Green Strategies and Naïve Farmers **(b)** Green Strategies and Strategic Farmers

Figure A4.3: Social Welfare Loss(%) for Political Allocation to Farmers = 70%

A5 Heavy–tailed draws for environmental values

As a final robustness check we relax the assumption that environmental and agricultural values are uniformly distributed on $[0.1, 10]$. Instead, we generate environmental values e from a Pareto distribution with shape parameter $\alpha_e \in [1.5, 2.0]$, rescaled to $[0.1, 10]$. This produces a heavy–tailed distribution with a small number of extremely high–value plots. Agricultural values a are kept on the baseline uniform $[0.1, 10]$ scale, so that the heavy tails apply only to e. The joint ranks of e and a are correlated via a t–copula with degrees of freedom $\nu = 4$, allowing for tail dependence. Formally, if $U \sim \text{Uniform}(0, 1)$ then

$$e = F^{-1}_{\text{Pareto},\alpha_e}(U), \quad a \sim U[0.1, 10],$$

with e and a linked through the copula. We then rescale e to the $[0.1, 10]$ support for comparability with the baseline.

Results are displayed in figures A5.1 and A5.2. Heavy–tailed draws make conservation outcomes hinge on a few "superstar" plots, increasing sampling variability across replications and hence widening confidence intervals. Against naïve developers, Max–Environment benefits by consistently targeting the superstars that developers ignore. Against strategic developers, both Max–Environment and Hot–Spot converge on the same superstars, narrowing the gap between them. Despite these differences in level and dispersion, the leakage–dependent comparative statics and the overall ranking of strategies remain unchanged.

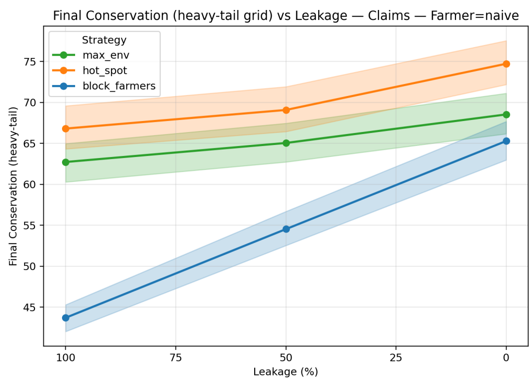

(a) Final Conservation with heavy tailed e-values

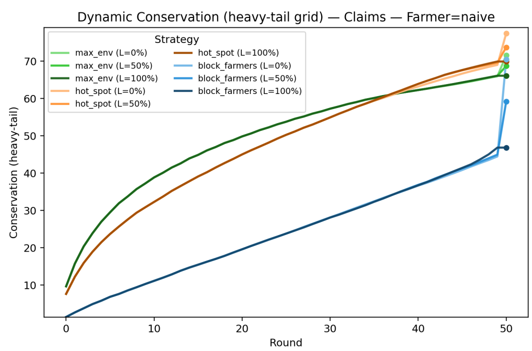

(b) Dynamic trajectories with heavy tailed e-values

Figure A5.1: Conservation with Heavy Tailed e Values: Green Strategies and Naïve Farmers

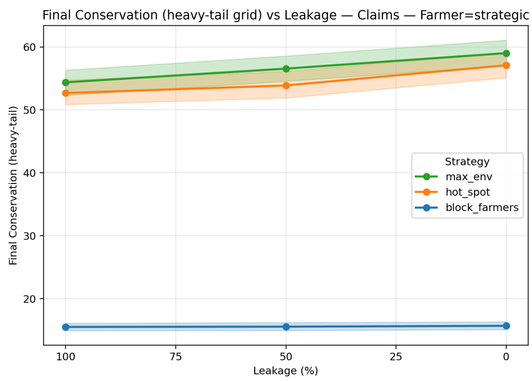

(a) Final Conservation with heavy tailed e-values

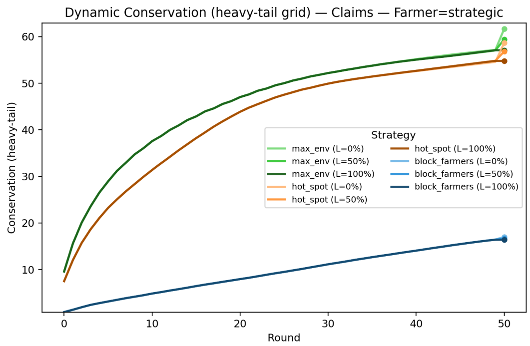

(b) Dynamic trajectories with heavy tailed e-values

Figure A5.2: Conservation with Heavy Tailed e Values: Green Strategies and Strategic Farmers

A6 Spatial robustness: S1 version

As a further robustness exercise, we introduce a simple spatial spillovers framework into the Claims World that we label S1. At the start of each simulation we randomly pair plots, but the links remain hidden until one member of a pair is claimed. When a plot is claimed by the Greens, its partner's environmental value is increased by a fixed amount (a "connectivity bonus"); when a plot is claimed by the Farmers, its partner's environmental value is decreased by the same amount (a "fragmentation penalty"). These adjustments change the effective rank ordering of the remaining available plots and therefore feed back into subsequent choices.

We implement the exercise with $m = 50$ disjoint pairs (i.e. all 100 plots linked) and bump/penalty magnitudes up to 2.0 on the $[0.1, 10.0]$ value scale. Results for both static final outcomes and dynamic trajectories are displayed in figures A6.1 and A6.2. As expected, this shifts realized conservation levels because effective values are altered during play, generally increasing final conservation scores across all Green strategies. However, across leakage values the qualitative patterns of performance by strategy—both in static outcomes and in dynamic trajectories—remain visually very similar to the baseline. We therefore interpret this exercise as a strong stress test that introduces within-game path dependence without overturning our main results.

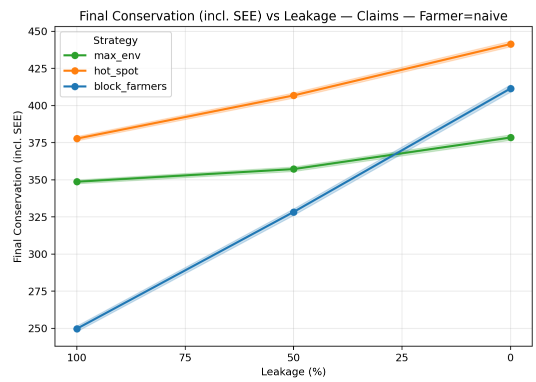

(a) Final Conservation with spatial effects

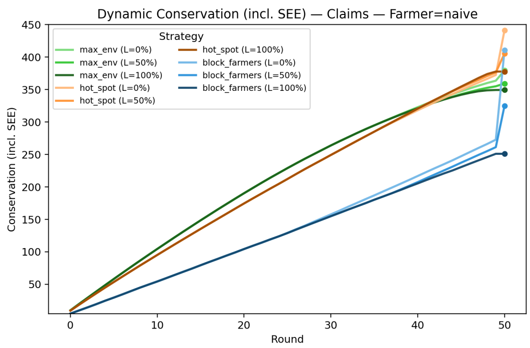

(b) Dynamic trajectories with spatial effects

Figure A6.1: Conservation with Spatially Linked Plots: Green Strategies and Naïve Farmers

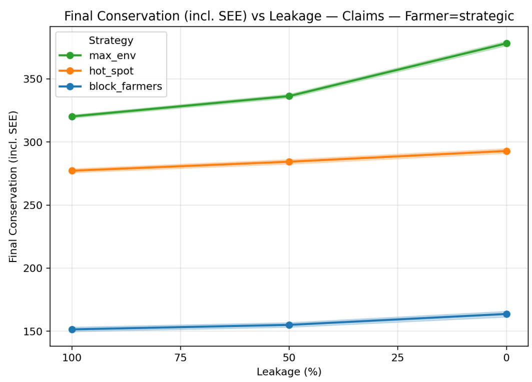

(a) Final Conservation with spatial effects

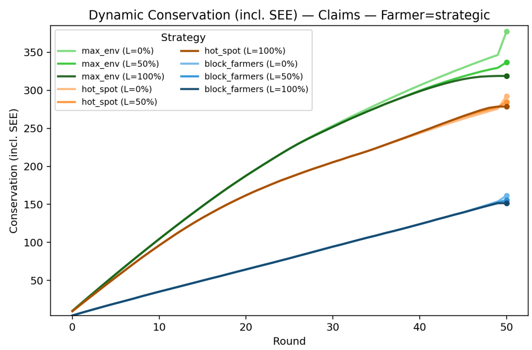

(b) Dynamic trajectories with spatial effects

Figure A6.2: Conservation with Spatially Linked Plots: Green Strategies and Strategic Farmers

A6.1 Spatial Externalities on a Graph (S2): Setup, Equivalence to S1, and Why It Favors Max–Environment

This appendix formalizes a graph–based spatial externality variant (“S2”) that allows spillovers to accumulate over multiple neighbors. We (i) define the environment; (ii) show that the S1 “paired shocks” case is a special case of S2; and (iii) establish that under mild assumptions S2 systematically advantages value–first conservation (Max–Environment) over threat–weighted rules by creating complementarities from clustering. Throughout, we keep the *Claims World* timing and accounting from the main text—including the PSE/DLE decomposition that credits residual plots at the end to visually separate displacement from direct strategy effects.

A6.1.1 S2: Graph–based spillovers

Let $P = \{1, \dots, N\}$ denote plots. Let $W = (W_{ij})$ be a nonnegative, symmetric adjacency matrix on P with $W_{ii} = 0$. For transparency we take W to be row–normalized (each row sums to k_i or to 1; both conventions are covered below).

The *effective* environmental value of plot i after some claims have been made is

$$e_i^{\text{eff}}(S_g, S_f) \;=\; e_i \;+\; \eta_G \sum_{j \in S_g} W_{ij} \;-\; \eta_F \sum_{j \in S_f} W_{ij}, \tag{1}$$

with $\eta_G, \eta_F \geq 0$. The Green team’s purchased/claimed component of conservation (our PSE°) is the sum of base values on Green–claimed plots,

$$\text{PSE}^{\circ}(S_g) \;=\; \sum_{i \in S_g} e_i,$$

and we define the *Spatial Externality Effect* (SEE) on the same claimed set as

$$\text{SEE}(S_g, S_f) \;=\; \sum_{i \in S_g} \{e_i^{\text{eff}}(S_g, S_f) - e_i\} \;=\; \eta_G \sum_{i \in S_g} \sum_{j \in S_g} W_{ij} \;-\; \eta_F \sum_{i \in S_g} \sum_{j \in S_f} W_{ij}.$$

As in the body of the paper, the *Displacement–Leakage Effect* (DLE) credits the residual unclaimed plots at the end (mechanical, leakage–driven). Hence the exact identity

$$\text{(2)}\quad \text{Final Conservation} \;=\; \underbrace{\mathrm{PSE}^{\circ}(S_g)}_{\text{strategy on base } e} \;+\; \underbrace{\mathrm{SEE}(S_g, S_f)}_{\text{spatial spillovers on claimed}} \;+\; \underbrace{\mathrm{DLE}(\text{Residual})}_{\text{mechanical residual credit}},$$

which mirrors our PSE/DLE convention while making spatial gains visible and additive.[1]

A6.1.2 S1 $\subset$ S2: Equivalence under a perfect matching

Lemma A6.1 (S1 is the k=1 special case of S2)**.** *Suppose W encodes a disjoint union of pairs: for each i there exists a unique $m(i)$ with $W_{i,m(i)} = W_{m(i),i} = 1$ and $W_{ij} = 0$ otherwise (no self–loops; no other edges). If B1 applies a λ_G boost to a partner's e when the first of a pair is claimed by Greens and a λ_F penalty when claimed by Farmers, then (1) coincides with B1 upon setting $\eta_G = \lambda_G$ and $\eta_F = \lambda_F$ (for unnormalized W). Under row–normalization the same equivalence holds since degrees are $k_i = 1$ for all i.*

Proof. With disjoint pairs, $\sum_{j\in S_g} W_{ij}$ (resp. S_f) equals 1 iff $m(i) \in S_g$ (resp. S_f); otherwise 0. Substituting into (1) reproduces the S1 update rule. ☐

Why S2 favors clustering and, therefore, Max–Environment

Define the set function on Green claims

$$\text{(3)}\qquad F(S_g; S_f) \;=\; \sum_{i\in S_g} e_i \;+\; \eta_G \sum_{i\in S_g}\sum_{j\in S_g} W_{ij} \;-\; \eta_F \sum_{i\in S_g}\sum_{j\in S_f} W_{ij},$$

so that $\mathrm{PSE}^{\circ} + \mathrm{SEE} = F(S_g; S_f)$.

Proposition A6.2 (Monotonicity and supermodularity)**.** *If $W_{ij} \geq 0$ and $\eta_G \geq 0$, $F(\cdot\,; S_f)$ is (i) monotone in S_g and (ii)* supermodular *in S_g (increasing differences). Equivalently, the marginal gain from adding i,*

$$\Delta_i(S_g; S_f) \;=\; e_i \;+\; \eta_G \sum_{j\in S_g} W_{ij} \;-\; \eta_F \sum_{j\in S_f} W_{ij},$$

[1] We keep residual crediting at the final round so that PSE/SEE dynamics remain visually separable from DLE, exactly as in the main text.

is (weakly) larger when evaluated at larger Green clusters S_g. In particular, with $\eta_F = 0$ (or holding S_f fixed) the gain from adding a plot rises with the mass of nearby Green claims.

Proof. Monotonicity is immediate since all terms added by enlarging S_g are nonnegative when $\eta_G \geq 0$. Supermodularity follows because the cross term $\sum_i \sum_j W_{ij} \mathbf{1}\{i \in S_g\} \mathbf{1}\{j \in S_g\}$ is quadratic with positive cross–partials in the Green indicators; e enters additively. $\square$

The proposition formalizes a *connectivity dividend*: as a Green cluster grows, nearby plots become incrementally more attractive, *independently of agricultural values a*. This has two immediate implications for the strategy ranking considered in the paper:

Corollary A6.3 (B2 tilts the playing field toward Max–Environment)**.** *Assume (A1) W is exogenous and independent of (e, a); (A2) $\mathrm{corr}(e, a) = 0$ (as in the baseline);*[3] *(A3) leakage is not zero (so displacement is imperfect, consistent with Section 2.2.4). Then, among the fixed heuristics studied in the body of the paper, S2 strictly increases the relative advantage of* MAX–ENVIRONMENT *over* HOT SPOT *and* BLOCK FARMERS *in expected* $\mathrm{PSE}^{\circ} + \mathrm{SEE}$*, and the gap grows in η_G and in average degree $\bar{k}$ of W.*

Proof sketch. The myopic one–step gain under S2 is $\Delta_i(S_g; S_f) = e_i + \eta_G \sum_{j \in S_g} W_{ij} - \eta_F \sum_{j \in S_f} W_{ij}$, which *does not* depend on a_i. Under (A1)–(A2), any rule that adds a positive weight on a_i (e.g. HOT SPOT or BLOCK FARMERS) misranks alternatives relative to Δ_i with probability one, lowering the expected one–step gain. Supermodularity then amplifies early misrankings: when a heuristic is pulled away from high–e seeds, subsequent $\sum_{j \in S_g} W_{ij}$ terms are smaller, compounding the shortfall. By contrast, MAX–ENVIRONMENT is aligned with the e_i part of Δ_i and (via supermodularity) naturally "rolls" along the Green frontier, harvesting the connectivity dividend. Increasing η_G or $\bar{k}$ raises the slope of these complementarities, magnifying the relative advantage. $\square$

Remarks. (i) When leakage tends to one ($L \to 1$), the BAU displacement channel shuts down and the Green best response in the non–spatial game reduces to a preempt–or–value rule (Appendix A6), which already disfavors pure threat–chasing; S2's spatial complementarity *further* shifts weight toward the value component. (ii) If the links in W are hidden from

[3] The paper reports the $\rho \neq 0$ cases separately; the same logic applies while the a weight ceases to be pure noise.

players ("unobserved" S2), SEE remains an ex post accounting gain but does not influence choices *ex ante*; in that knife–edge case S2 behaves more like S1 with random pairing.[4]

Comparative statics and calibration to S1

To compare S1 and S2 on a common scale, let $\bar{k}$ denote the average degree of W. If S1 uses shocks λ_G, λ_F on disjoint pairs, then a convenient calibration for row–normalized S2 is

$$\eta_G \approx \bar{k}\,\lambda_G, \qquad \eta_F \approx \bar{k}\,\lambda_F,$$

so that the *per–neighbor* increment in e^{eff} matches the B1 shock. Under this mapping,

- S1 with M disjoint pairs corresponds to S2 on a graph where each plot has degree $k \in \{0, 1\}$ (a partial matching).
- As degree increases ($k \geq 2$), S2 creates multi–neighbor accumulation that S1 cannot replicate. The connectivity dividend (and thus the tilt toward value–first play) rises in η_G and in $\bar{k}$.

Take–away for the reader

S2 embeds a simple economic force: clustering by Greens *creates* value on neighbors (supermodularity), whereas weighting by agricultural threat does not. Because the one–step gain Δ_i is orthogonal to a_i under (A1)–(A2), heuristics that put positive weight on a_i are systematically misaligned with the spatial objective and underperform relative to Max–Environment—and the shortfall grows with the strength and scope of spillovers. This is precisely why we use S1 (dyadic spillovers) for robustness: it is a *conservative* stress that is less favorable to Max–Environment than the more realistic multi–neighbor S2.

[4] See Section 2 for the non–spatial nature of the grid in the baseline setup; S2 is introduced purely as a robustness lens, not as a claim about geographic adjacency in the main simulations.

A7 Budget World: Formal Set Up

A7.1 Budgets

Let $\{(e_i, a_i)\}_{i=1}^N$ denote the environmental and agricultural values of $N = n^2$ plots. The price of each plot equals its agricultural value: $\text{Price}(p_i) = a_i$.

Farmer and Green budgets are calibrated as shares of the total agricultural value:

$$B_f^0 = s \sum_{i=1}^N a_i, \quad B_g^0 = (1-s) \sum_{i=1}^N a_i,$$

where $s \in (0, 1)$ is the Farmer's share. In the baseline case $s = 0.5$, the two sides together can purchase the entire grid under full leakage.

Farmer Strategies

Farmer strategies mirror the Claims World:

1. **Naïve budgeter:** purchase the highest-a_i plots until budget is exhausted.
2. **Strategic budgeter:** identify "risky" plots (those targeted under a Green Max Environment strategy with B_g^0) and prioritize them in descending a_i, then fill remaining budget with high-a_i "safe" plots.

A7.2 Green Strategies

Greens may adopt four strategies:

1. **Max Environment:** purchase the affordable plot with highest e_i.
2. **Hot Spot:** purchase the affordable plot maximizing $e_i \cdot a_i$.
3. **Block Farmers:** purchase the affordable plot with highest a_i.

4. **Max Efficiency (ratio-greedy):** purchase the affordable plot maximizing $r_i = e_i/a_i$. This is the standard cost-effectiveness heuristic widely used in conservation planning (a greedy solution to the 0–1 knapsack problem).[5]

A7.3 Gameplay and Leakage

Teams alternate purchases until neither can afford additional plots. If Greens capture a Farmer-BAU plot p_i, the Farmer's budget is reduced by

$$B_f \leftarrow B_f - (1 - L)a_i,$$

where $L \in [0, 1]$ denotes leakage. At $L = 1$ no budget is burned; at $L = 0$ the full value of the blocked plot is lost.

Outcome Metrics

We distinguish between direct purchases and residual gains:

- **Purchased Conservation (PC):** $PC = \sum_{p_i \in S_g} e_i$.
- **Displacement–Leakage Effect (DLE):** $DLE = \sum_{p_i \in \text{Residual}} e_i$, credited only in budget-parity cases ($s = 0.5$) when residuals reflect displaced Farmer BAU plots.
- **Final Conservation:** $C_{\text{final}} = PC + DLE$.
- **Welfare Loss:** decline from the maximum attainable welfare, reported separately for purchased outcomes and (under parity) final outcomes including DLE.
- **Event Additionality:** increases by e_i when Greens purchase a Farmer-BAU plot, and decreases by e_i when Farmers purchase a Green-BAU plot.

[5] In the fractional relaxation of the knapsack, the optimal solution is to select all items with $e_i/a_i > \lambda^*$ for some threshold λ^*. The ratio-greedy rule mirrors this structure while remaining heuristic in the 0–1 case.

A8 Theory Addendum: Equilibrium Structure and Budget-World Dominance

This addendum connects the simulation patterns in the baseline *Claims World* (Section 2) and the *Budget World* (Section 4) to tractable best–response logic in polar cases and to exact tiny–grid equilibria that benchmark our heuristics against optimizing developers. Closed–form equilibria remain elusive due to state dimensionality; our contribution is to establish transparent limiting cases, show how they rationalize the simulation rankings, and verify them on small grids. Code and results are available from the authors.

A8.1 Claims World: polar leakage cases

Full leakage ($L = 1$): a myopic preempt–or–value rule.

Proposition A8.1 (Claims World, $L = 1$). *Suppose leakage is full ($L = 1$) and Farmers play the naïve Profit Maximizer (Section 2.2.2). Then: (i) the Green value function satisfies*

$$V_g(A) = \max_{i \in A} \Big\{ e_i + V_g(A \setminus \{i,\, j(A \setminus \{i\})\}) \Big\},$$

where $j(s) = \arg\max_{i \in s} a_i$ is the Farmer's threatened plot; (ii) a Green best reply is preempt–or–value*: pick $j(s)$ if $e_{j(s)} > E_{safe}(s)$, else pick the highest–e safe plot; and (iii) (naïve Farmer, preempt–or–value) is a Markov–perfect equilibrium.*

Remark A8.2 (Strategy ranking as $L \uparrow 1$). With $\mathrm{corr}(e, a) = 0$, ranking by a is independent of e, so pure threat–chasing (Block Farmers) samples e almost at random and is dominated by Max Environment on PSE, consistent with our figures and tables.

Zero leakage ($L = 0$): a threshold rule with a BAU bonus.

Proposition A8.3 (Claims World, $L = 0$). *With $L = 0$ and naïve Farmers, a Green best reply at state s selects*

$$i^\star \in \arg\max_{i \in A(s)} \Big\{ e_i + \theta_s \cdot \mathbf{1}\{\, i \in S_f^{BAU}(s) \,\} \Big\},$$

where θ_s is the expected e on the marginal residual plot. In words: protect by e, with a BAU "bonus" θ_s for threatened plots.

Corollary A8.4 (Intermediate leakage)**.** *For $L \in (0, 1)$ the BAU bonus scales as $(1 - L)\theta_s$, so the hotspot tilt shrinks smoothly to zero as $L \uparrow 1$.*

Lemma A8.5 (Strategic Farmers)**.** *If Farmer BAU plots are stochastically larger in e (our Strategic Farmer), the threshold rule collapses to MAX ENVIRONMENT: the BAU indicator and e move together, so the bonus does not overturn e–maximization.*

A8.2 Budget World: efficiency rules and dominance

We study Purchased Conservation (PC), which strips out the residual-crediting channel through which leakage primarily affects total conservation. In the full dynamic game the purchased set can still vary with L via Farmer budget burn, but under thick budgets this feedback is second-order; the selection effect that drives the Max Environment > Max Efficiency comparison is therefore essentially independent of L. Farmer removal hazards are assumed weakly increasing in a signal strictly increasing in a (covering both naïve rank–by–a and Strategic high–a targeting). Items are i.i.d. (e_i, a_i) with continuous marginals; unless stated we assume a nonnegative association structure ($\mathrm{MTP}_2/\mathrm{TP}_2$).

A8.2.1 General dominance on Purchased Conservation.

Simulation results (Section 4.1) show that MAX ENVIRONMENT systematically outperforms MAX EFFICIENCY on purchased conservation (PC) across all leakage levels and against both naïve and Strategic Farmers. Figure A8.1 zooms in on the head-to-head comparison between these two rules, plotting PC against leakage with 95% confidence bands. The dominance of MAX ENVIRONMENT is visually clear and essentially leakage-invariant. Table A8.1 reports win-rates and PC gaps across 500 replications per leakage level for each Farmer type, corroborating the near-universal dominance observed in the plots.

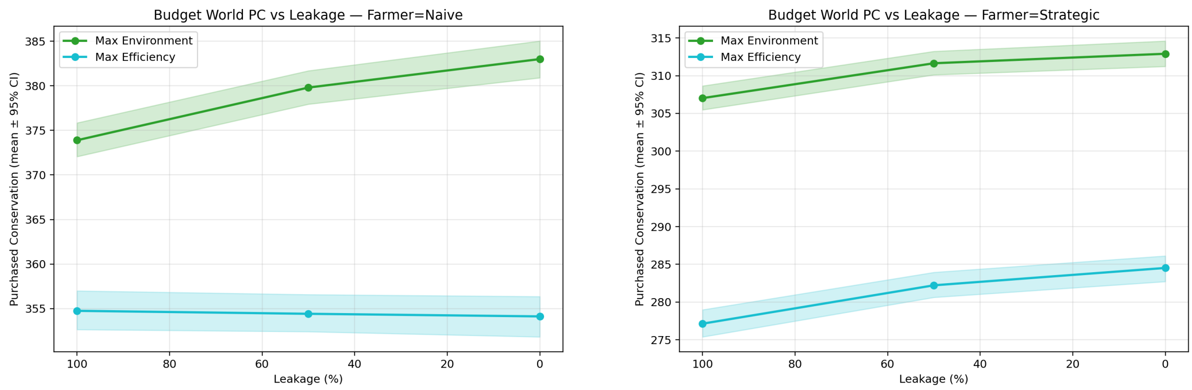

(a) naïve Farmer **(b)** Strategic Farmer

Figure A8.1: Budget World head-to-head Purchased Conservation (PC) vs. leakage for naïve and Strategic Farmers. Max Environment (green) strictly dominates Max Efficiency (teal) in essentially all replications; curves are nearly parallel across leakage, indicating a selection effect that is independent of L. Shaded bands: ±95% across 500 replications.

Table A8.1: Budget World: Max Environment vs. Max Efficiency on PC

Farmer	L	Wins	Win-rate	Mean gap	P10	P90
naïve	0.0	500/500	1.000	28.9	16.4	42.2
	0.5	500/500	1.000	25.4	12.5	38.4
	1.0	498/500	0.996	19.1	7.7	30.5
Strategic	0.0	500/500	1.000	28.4	18.7	39.2
	0.5	500/500	1.000	29.4	18.4	41.3
	1.0	500/500	1.000	29.9	19.1	40.7

PC gap = PC(MaxEnv)−PC(MaxEff). P10/P90 denote the 10th/90th percentiles of the PC gap distribution across replications. Entries are computed from 500 replications at each leakage level. See Theorem A8.9 for the formal statement.

A8.2.2 Constructive counterexample.

Proposition A8.6 (Counterexample in Budget World)**.** *Under budget parity, there exist instances and leakage levels L for which* Max Efficiency *(ratio–greedy e/a) yields strictly lower final conservation than* Max Environment *(value–greedy e), hence is not a best response.*

Constructive example. Consider three plots with (e, a) values $(9, 9), (8, 8), (7, 1)$, parity budgets $B_g^0 = B_f^0 = 9$, and naïve Farmers. Ratios are $1, 1, 7$. If Greens play Max Efficiency, they take $(7, 1)$, the Farmer buys $(9, 9)$, and Greens then get $(8, 8)$: final conservation $= 15$. If Greens play Max Environment, they first take $(9, 9)$, burning $(1 - L) \cdot 9$ of Farmer budget; the Farmer's remaining budget is $B'_f = 9L$, so for $L < 8/9$ the Farmer cannot afford $(8, 8)$, which Greens then secure, yielding final conservation $= 17 > 15$. □

Remark A8.7. This mechanism extends whenever a high–(e, a) hotspot lies in Farmer BAU: ratio–greedy avoids it on cost grounds and forfeits both the hotspot and, for $L < 1$, the budget–burn benefit.

A8.2.3 Theory: one–step lemma and asymptotic dominance.

Lemma A8.8 (One–step PC dominance of e–greedy)**.** *Fix any state in Budget World. Let $k = \arg\max e$ among the affordable plots and let a Green rule G (possibly history–dependent) select $j \neq k$ with $e_j < e_k$. Under nonnegative E–A association (MTP$_2$/TP$_2$ suffices) and a Farmer removal hazard weakly increasing in an a–signal, the expected shortfall in PC at the next Green selection time from deferring k is strictly positive:*

$$\mathbb{E}[\Delta\mathrm{PC} \mid \textit{defer } k] = \mathbb{E}[\mathbf{1}\{k \textit{ is removed before } G \textit{ returns}\} \cdot e_k] + \mathbb{E}[e_k - e_j] \;>\; 0.$$

This one–step comparison is independent of residual crediting*: since PC excludes the DLE term, the conditional shortfall above is identical for all $L \in [0, 1]$* given the state*. Leakage can affect PC only indirectly by shifting the future distribution of states (budgets and hence opportunity sets), which is addressed in the global result under thick budgets.*

Theorem A8.9 (Asymptotic PC dominance under nonnegative association (pointwise in leakage)). *Consider Budget World with budget parity. Let* (e_i, a_i) *be i.i.d. with continuous marginals and* MTP_2 *dependence. Suppose* $\max_i a_i/B_g \to 0$ *as* $N \to \infty$ *("thick budgets"), and Farmer hazards are weakly increasing in the a–signal. Then, for any* $L \in [0,1]$*,*

$$\liminf_{N\to\infty} \Big\{ \mathbb{E}[\mathrm{PC}(\textsc{MaxEnv})] - \mathbb{E}[\mathrm{PC}(G)] \Big\} \;\geq\; 0$$

for any alternative Green rule G*, with strict inequality whenever* G *defers the current max–e item with positive probability. The dominance conclusion is* pointwise *in leakage: it holds for each fixed* $L \in [0,1]$*. Moreover, since PC excludes the residual DLE crediting term, the* statewise *one–step comparison underlying the proof is the same for all L; any dependence of PC on L operates only through leakage-induced shifts in the state path, which are asymptotically negligible under thick budgets.*

Sketch. Sum the one–step gains from Lemma A8.8 along the play. The only correction is a *budget–feedback* remainder from $a_k \neq a_j$ at disagreements (future affordability sets differ). Under $\max_i a_i/B_g \to 0$, this remainder is $o_p(1)$, so expected cumulative one–step gains dominate in the limit. Since PC excludes the residual DLE crediting term, L does not enter the conditional one–step gain; any dependence of PC on L operates only through leakage-induced shifts in the state path, which vanish asymptotically under thick budgets. □

Remark A8.10 (Finite grids and budget feedback). In small or tight–budget instances a single very high–e, very high–a purchase can crowd out many medium–e purchases, so local gains need not globalize at finite N. Exact tiny–grid equilibria show this knife–edge: on 2×2 the mean PC gap $\mathbb{E}[\mathrm{PC}(\textsc{MaxEnv}) - \mathrm{PC}(\textsc{MaxEff})]$ is negative across ρ; on 3×3 it is positive and grows with ρ (see Subsection A8.3 and Table A8.3).

A8.2.4 Comparative statics in the dependence parameter.

Proposition A8.11 (Monotone comparative statics). *Let* (E, A) *be generated by a one–parameter family* $\{C_\theta\}_{\theta\in[-1,1]}$ *(e.g. a Gaussian copula) that is increasing in the supermodular order. Define*

$$\Delta(\theta) = \mathbb{E}[\mathrm{PC}(\textsc{MaxEnv}) - \mathrm{PC}(\textsc{MaxEff})]\,.$$

Under the assumptions of Theorem A8.9, $\Delta(\theta)$ is (weakly) increasing in θ, with $\lim_{\theta\to-1}\Delta(\theta)=0$ and $\Delta(\theta)\geq 0$ for $\theta\geq 0$ asymptotically.

Empirical alignment. In our 10×10 simulations the PC gap is large and essentially flat in L (leakage–invariant); in tiny grids the mean gaps match the comparative statics: negative on 2×2, positive and growing with ρ on 3×3.

A8.3 Tiny–grid equilibrium enumeration

We compute exact equilibria on 2×2 and 3×3 in Budget World (PC only) and on 2×2, 3×3, 4×4 in Claims World, benchmarking heuristics against the equilibrium Green strategy against Farmers playing their equilibrium optimizing strategy. The results in Tables A8.2 and A8.3 confirm that the strategy rankings observed in the Monte Carlo simulations are structural rather than heuristic artifacts. In *Claims World*, equilibrium Green play behaves as an e–first rule with a discrete BAU bonus: empirically, the Hot Spot heuristic ($e\times a$) most closely matches equilibrium outcomes, while Max Environment, Block Farmers, and ratio–greedy rules diverge with grid size. In *Budget World*, where agricultural value a_i also determines price, equilibrium behavior converges toward the value–greedy (Max Environment) rule.

Across both Worlds, first–move effects diminish rapidly with dimension, and the dominance of value–first play on Purchased Conservation is monotone in the e–a correlation ρ. These small-grid equilibria therefore validate the broader simulation results and clarify the transition from *adversarial siting* (where threat-weighting aids performance) to *adversarial procurement* (where cost weighting becomes detrimental).

Table A8.2: Tiny–grid equilibrium benchmarks (Claims World, by correlation ρ)

Grid	ρ	MaxEnv	Hot Spot	Block Farmers	Ratio	First–move agree
2×2	-0.8	421/2000 (0.30 ± 0.80)	1103/2000 (1.76 ± 2.40)	1953/2000 (5.53 ± 2.81)	668/2000 (0.58 ± 1.14)	912/2000 (45.6%)
	-0.4	406/2000 (0.36 ± 0.96)	850/2000 (0.94 ± 2.11)	1760/2000 (4.31 ± 3.52)	1024/2000 (1.32 ± 1.86)	1052/2000 (52.6%)
	0.0	342/2000 (0.34 ± 1.02)	775/2000 (0.60 ± 1.77)	1606/2000 (3.29 ± 3.42)	1165/2000 (1.82 ± 2.24)	1141/2000 (57.0%)
	$+0.4$	265/2000 (0.25 ± 0.84)	635/2000 (0.28 ± 1.35)	1364/2000 (2.18 ± 2.97)	1401/2000 (2.42 ± 2.47)	1317/2000 (65.8%)
	$+0.8$	117/2000 (0.08 ± 0.45)	441/2000 (0.12 ± 0.77)	943/2000 (0.85 ± 1.64)	1598/2000 (2.79 ± 2.49)	1599/2000 (80.0%)
3×3	-0.8	1331/2000 (1.62 ± 1.87)	1138/2000 (1.25 ± 2.10)	1990/2000 (9.47 ± 3.75)	1386/2000 (1.80 ± 2.01)	278/2000 (13.9%)
	-0.4	1524/2000 (2.88 ± 2.61)	1005/2000 (0.71 ± 2.06)	1944/2000 (8.16 ± 4.61)	1732/2000 (4.38 ± 3.21)	415/2000 (20.8%)
	0.0	1501/2000 (2.89 ± 2.65)	958/2000 (0.44 ± 2.02)	1865/2000 (6.73 ± 4.83)	1897/2000 (6.07 ± 3.38)	607/2000 (30.3%)
	$+0.4$	1380/2000 (2.25 ± 2.44)	864/2000 (0.20 ± 1.87)	1748/2000 (4.89 ± 4.42)	1957/2000 (7.29 ± 3.36)	950/2000 (47.5%)
	$+0.8$	931/2000 (0.92 ± 1.52)	825/2000 (0.07 ± 1.20)	1620/2000 (2.37 ± 2.89)	1991/2000 (7.66 ± 3.06)	1333/2000 (66.6%)
4×4	-0.8	1757/2000 (3.17 ± 2.45)	1613/2000 (3.16 ± 3.30)	2000/2000 (20.01 ± 4.84)	1770/2000 (3.34 ± 2.59)	118/2000 (5.9%)
	-0.4	1905/2000 (5.79 ± 3.22)	1243/2000 (1.15 ± 3.08)	1997/2000 (16.77 ± 6.66)	1956/2000 (8.24 ± 4.11)	273/2000 (13.7%)
	0.0	1894/2000 (5.61 ± 3.12)	994/2000 (0.05 ± 2.96)	1960/2000 (13.24 ± 6.71)	1988/2000 (11.45 ± 4.43)	494/2000 (24.7%)
	$+0.4$	1868/2000 (4.61 ± 2.86)	869/2000 (−0.39 ± 2.52)	1881/2000 (9.63 ± 6.18)	1999/2000 (14.05 ± 4.10)	801/2000 (40.1%)
	$+0.8$	1671/2000 (2.28 ± 2.00)	921/2000 (−0.24 ± 1.77)	1784/2000 (4.85 ± 4.07)	2000/2000 (15.58 ± 3.60)	1236/2000 (61.8%)

Cells report *wins/trials* (EQ beats the heuristic playing against the equilibrium Farmer strategy) stacked over the mean Green payoff gap (EQ − heuristic) with its s.d. "Ratio" is the ratio–greedy rule (e/a). Reps: 2,000 per ρ (seed 321).

Table A8.3: Tiny–grid equilibrium benchmarks (Budget World, by correlation ρ)

Grid	ρ	EQ vs MaxEnv	EQ vs MaxEff	EQ vs Hot Spot	EQ vs Block Farmers	MaxEnv vs MaxEff
	-0.8	389/2000	443/2000	1483/2000	1892/2000	271/2000
		(1.49 ± 3.84)	(0.63 ± 1.78)	(9.49 ± 8.35)	(18.49 ± 9.09)	(-0.86 ± 3.82)
2×2	-0.4	672/2000	616/2000	1435/2000	1813/2000	361/2000
		(2.58 ± 4.59)	(1.17 ± 2.58)	(8.11 ± 7.59)	(15.50 ± 9.34)	(-1.42 ± 5.05)
	0.0	857/2000	746/2000	1413/2000	1727/2000	418/2000
		(2.86 ± 4.62)	(1.48 ± 2.96)	(6.76 ± 6.85)	(12.72 ± 9.29)	(-1.38 ± 5.36)
	$+0.4$	1009/2000	869/2000	1379/2000	1587/2000	510/2000
		(3.05 ± 4.38)	(1.90 ± 3.29)	(5.58 ± 6.00)	(9.22 ± 8.17)	(-1.16 ± 5.39)
	$+0.8$	1134/2000	1044/2000	1308/2000	1417/2000	638/2000
		(3.02 ± 4.04)	(2.18 ± 3.24)	(4.04 ± 4.65)	(5.40 ± 5.66)	(-0.85 ± 4.99)
	-0.8	1304/2000	1489/2000	1831/2000	2000/2000	712/2000
		(5.52 ± 6.96)	(6.81 ± 7.68)	(17.73 ± 13.34)	(39.15 ± 13.53)	(1.29 ± 6.12)
3×3	-0.4	1410/2000	1630/2000	1810/2000	1998/2000	897/2000
		(5.63 ± 6.43)	(7.10 ± 6.85)	(15.14 ± 11.72)	(33.65 ± 12.92)	(1.46 ± 6.91)
	0.0	1443/2000	1690/2000	1780/2000	1982/2000	1016/2000
		(5.65 ± 6.28)	(7.44 ± 6.69)	(12.72 ± 10.05)	(27.84 ± 12.03)	(1.79 ± 7.37)
	$+0.4$	1511/2000	1756/2000	1766/2000	1957/2000	1119/2000
		(5.55 ± 5.79)	(7.35 ± 6.07)	(10.27 ± 8.43)	(20.49 ± 10.99)	(1.80 ± 7.17)
	$+0.8$	1578/2000	1781/2000	1753/2000	1917/2000	1197/2000
		(4.43 ± 4.43)	(6.50 ± 5.22)	(6.81 ± 5.74)	(10.62 ± 6.94)	(2.07 ± 6.26)

Cells report *wins/trials* (number of replications in which the equilibrium (EQ) Green strategy outperforms the heuristic playing against the equilibrium Farmer strategy) stacked over the mean Purchased Conservation (PC) gap (EQ − heuristic) with its s.d. Positive values favor EQ. “Ratio” denotes the ratio–greedy rule (e/a). Reps: 2,000 per ρ (seed 321). PC is leakage–invariant.

A9 Supplement to Empirical Exercise on Bolivia

A9.1 A short history of protected areas in Bolivia

Bolivia provides an insightful case study of the dynamic and contested nature of conservation decisions. With an extraordinary level of biodiversity and significant pressure for economic development, the history of Bolivian conservation efforts illustrate well the trade-offs and strategic interactions that our simulated framework addresses. Protected Areas (PAs) in Bolivia span national, state, and municipal jurisdictions, varying widely in size, purpose, and effectiveness, but have in common a restriction on the type and extent of development that may take place. As illustrated in figure A9.1(b), since 1939 the total land area under Protected Area status has gradually increased from zero to 35.4 million hectares, 32% of the national territory[6]. Figure A9.1 reveals that PA expansion in Bolivia has proceeded at least 3 times faster than agropastoral expansion since 1985 (30 million hectares for conservation versus 8 million hectares for agropastoral expansion).

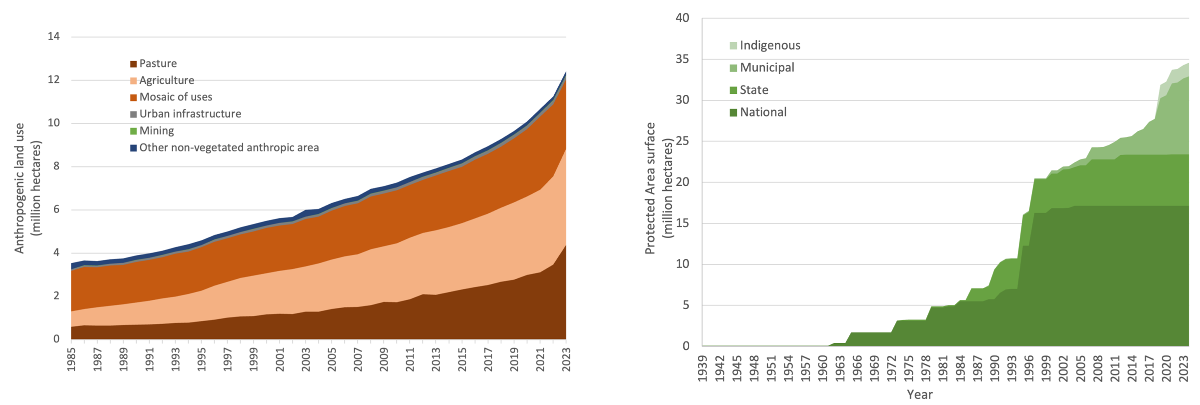

(a) Cumulative Anthropogenic Land Use

(b) Cumulative Protected Areas

Figure A9.1: Cumulative Land Use in Bolivia, 1930–2023

Formal conservation efforts began in 1939 with the creation of Sajama National Park to protect the endangered Polylepis forests around Bolivia's highest peak. Shortly thereafter, in 1942, the Tuni Condoriri National Park was established to safeguard critical water supplies for El Alto and La Paz. However, conservation activity remained sparse until the mid-1960s, when protected areas such as TIPNIS, Manuripi, and Eduardo Avaroa were established,

[6] Bolivia is among 35 countries in the world that have already reached the 30x30 goal of the CBD.

marking the first substantial extension of Bolivia's protected lands into regions with significant biodiversity value.

The late 1980s and 1990s constituted a "golden age" of Bolivian conservation, fueled largely by international funding through mechanisms like the pioneering 1987 Debt-for-Nature swap and extensive international support from NGOs. Major protected areas established during this period—including Amboró, Madidi, Carrasco, and Kaa-Iya del Gran Chaco—targeted regions with exceptionally high environmental values but often lower immediate agricultural threats, consistent with a strategy focused primarily on maximizing environmental benefits.

By contrast, the economic expansion from the early 2000s through around 2015 significantly reshaped conservation dynamics. Protected area establishment slowed markedly during this period due to greater economic incentives for agricultural and extractive development. At the same time, the new PAs that were established were often placed within or near the agricultural frontier, as is apparent in the Supplementary video A9.2 and as we illustrate below in figure 12, where we see that the agricultural potential of newly protected areas in the early 2000s is relatively high. The Laguna Concepción State-Level Wildlife Reserve, established in 2002, is a good example. Laguna Concepción is a natural lagoon designated a Ramsar Site, recognizing it as a wetland of international importance, but it is located in the center of Bolivia's Santa Cruz State near the heart of agricultural expansion in the country. Despite its Protected Area status, conservation of the Laguna has been a struggle, with significant agricultural and livestock encroachment by Mennonite colonies into the western and northern sectors, threatening the fragile ecosystem (Navia, 2022).

Following the economic downturn of the late 2010s and the subsequent pandemic crisis (2020–2023), Bolivia again intensified conservation efforts, rapidly designating new protected areas—particularly at the municipal level; the year 2019 saw the establishment of more than 4.2 million hectares of newly Protected Areas and the pandemic crisis of 2020–2023 coincided with the creation of 2.4 million hectares of primarily municipal protected areas.

Supplementary video A9.2 shows the dynamic expansion of agricultural land and Protected Areas from 1985 to 2024. Figures A9.3(a) and A9.3(b) show the average potential conservation and potential agricultural values of the land newly allocated to either agriculture or conservation (via Protected Area status), respectively, for each year from 1985 to 2023.

Figure A9.2: Supplementary Video: Dynamic expansion of Agricultural and Protected Area Land, 1985–2024
(Click the button to view the video).

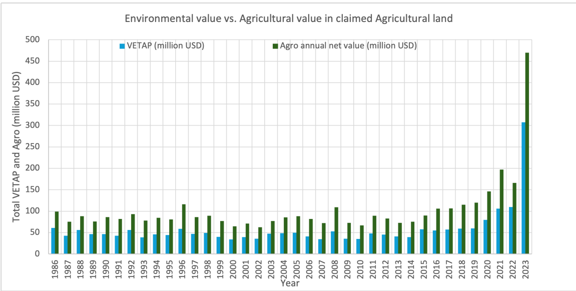

(a) New Agricultural Land, by year

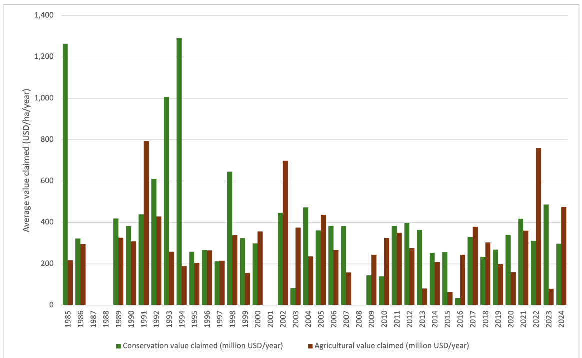

(b) New Protected Areas, by year

Figure A9.3: Potential Agricultural & Conservation Values of New Agricultural and Protected Land, (annual USD per ha)

A9.2 Data

In order to map the expansion of Protected Areas in Bolivia to our theoretical Green conservation strategies we combine annual (1985–2023) pixel-level land use data from Mapbiomas MapBiomas Bolivia Project (2024) and a detailed geo-referenced database of all Protected Areas in Bolivia, including type and year of establishment from SDSN Bolivia (SDSN Bolivia, 2025). A times series of Bolivian GDP per capita is from Our World in Data (Our World in Data, 2025).

In order to estimate the potential conservation values of land areas either designated as Protected Areas or put into agricultural production we take advantage of a high-resolution map of the value of ecosystem services from Andersen et al. (2024). The map, reproduced in figure A9.4(b), is expressed in *potential* annual dollar values per hectare (USD/ha/year) from the benefits of conservation, including provisioning values (timber, non-timber forest products, hunting and fishing, water), regulation values (biodiversity conservation, carbon sequestration, local climate regulation, pollination, and water treatment), and cultural values (tourism and recreation). Conservation values range from the lowest in the arid region of the southwestern

Bolivian Altiplano, to the highest in the dense Amazon jungle in the north and the biodiversity-rich mountainous valley regions between the highlands and the lowlands.

In order to estimate the Agricultural value of land we generate estimates of net agricultural potential per hectare, presented in figure A9.4(a), following the methodology from Andersen et al. (2023). Agricultural potential varies greatly across Bolivia due to differences in topography, soil quality, and climate, with the most profitable regions being those with climates appropriate for the production of high value crops, like fruits and berries[7]. To generate the map, we use detailed pixel-level geographical data on slope, soil quality, precipitation, and temperature distribution to develop a high-resolution Production Cost Factor. Combining this with information on the most common crop (or livestock), average yields, and prices in each municipality to generate pixel-level estimates of agricultural potential[8].

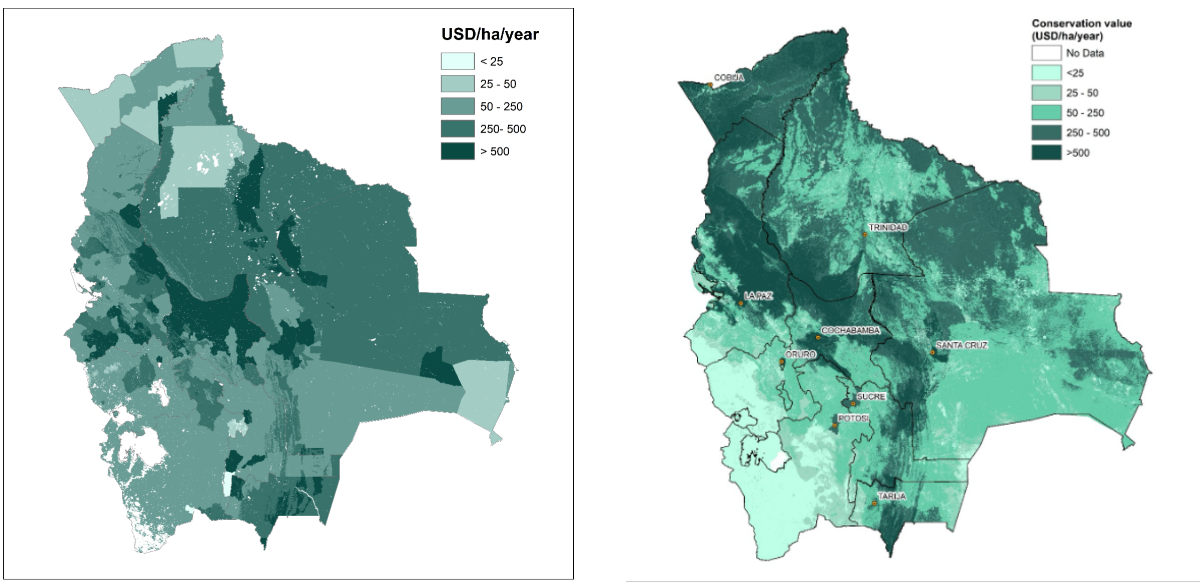

(a) Potential Agricultural Values

(b) Potential Conservation Values

Figure A9.4: Potential Agricultural & Conservation values (annual USD per ha)

[7] In practice, the private profitability of agriculture depends not only on agricultural potential but also on local infrastructure and proximity to markets, which are not reflected in these numbers. Nevertheless, the map provides an estimate of relative agricultural values.

[8] While the original map in Andersen et al. (2023) uses information on Protected Area status and infrastructure such as distance to roads to generate the Production Cost Factor estimates, for our purposes we have produced estimates that use only geo-physical conditions

A9.3 Aggregating pixel values to planning units

Let Bolivia be discretised into raster pixels indexed by $i = 1, \dots, I$. For each pixel i we observe:

- e_i: annual environmental value in USD per hectare per year (ecosystem services), and
- a_i: potential net annual agricultural value (proportional to USD per hectare per year).

Let planning units (PUs) be indexed by $j = 1, \dots, J$, and let $\mathcal{I}_j$ denote the set of pixels whose centres fall inside PU j.

Step 1: National pixel–level percentiles

We first construct national reference distributions for environmental and agricultural values across all land pixels. Let

$$\{e_i\}_{i=1}^{I}, \qquad \{a_i\}_{i=1}^{I}$$

denote the full sets of pixel–level values. Define the qth percentiles

$$P_q^{(e)} \quad \text{and} \quad P_q^{(a)} \qquad \text{for } q \in \{90, 95, 99\}$$

as the 90th, 95th and 99th percentiles of the distributions $\{e_i\}$ and $\{a_i\}$, respectively. These thresholds identify pixels in the upper tails of the national environmental and agricultural value distributions.

Step 2: PU–level means and tail fractions

For each PU j, we compute simple averages (the μ scores):

$$\mu_j^{(e)} = \frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} e_i, \qquad \mu_j^{(a)} = \frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} a_i, \tag{4}$$

which capture the typical environmental and agricultural value of a hectare in PU j.

To characterise how much of each PU lies in the upper tails of the national distributions, we define *tail fractions*. For environmental value:

(5) $$f_j^{(e)}(q) \;=\; \frac{\#\left\{i \in \mathcal{I}_j : e_i \geq P_q^{(e)}\right\}}{|\mathcal{I}_j|}, \qquad q \in \{90, 95, 99\},$$

and analogously for agricultural value:

(6) $$f_j^{(a)}(q) \;=\; \frac{\#\left\{i \in \mathcal{I}_j : a_i \geq P_q^{(a)}\right\}}{|\mathcal{I}_j|}, \qquad q \in \{90, 95, 99\}.$$

Thus $f_j^{(e)}(95)$, for example, is the fraction of PU j's area lying in the top 5% of environmental values nationally.

Step 3: Tail–Averaged Percentile Scores (TAPS)

To combine the information on *how high* and *how widespread* extreme values are within a PU, we construct simple Tail–Averaged Percentile Scores (TAPS). For environmental values:

(7) $$\mathrm{TAPS}_j^{(e)}(90) \;=\; 90{\cdot}f_j^{(e)}(90), \quad \mathrm{TAPS}_j^{(e)}(95) \;=\; 95{\cdot}f_j^{(e)}(95), \quad \mathrm{TAPS}_j^{(e)}(99) \;=\; 99{\cdot}f_j^{(e)}(99).$$

and similarly for agricultural values:

(8)

$$\mathrm{TAPS}_j^{(a)}(90) \;=\; 90 \cdot f_j^{(a)}(90), \quad \mathrm{TAPS}_j^{(a)}(95) \;=\; 95 \cdot f_j^{(a)}(95), \quad \mathrm{TAPS}_j^{(a)}(99) \;=\; 99 \cdot f_j^{(a)}(99).$$

These scores give more weight to pixels in higher percentiles and to PUs where such pixels occupy a larger share of the area.

Step 4: Percentile ranks across PUs

Because the raw means μ_j and TAPS are on different scales, we normalise them by converting each into a percentile rank across PUs. Let $F_\mu^{(e)}$ denote the empirical distribution of $\{\mu_j^{(e)}\}_{j=1}^{J}$. We define

(9) $$M_j^{(e)} \;=\; F_\mu^{(e)}(\mu_j^{(e)}),$$

so that $M_j^{(e)} \in [0,1]$ is the percentile rank of PU j's mean environmental value among all PUs. We similarly define percentile ranks for the environmental TAPS:

$$T_j^{(e,90)},\ T_j^{(e,95)},\ T_j^{(e,99)} \tag{10}$$

as the percentile ranks of $\mathrm{TAPS}_j^{(e)}(90)$, $\mathrm{TAPS}_j^{(e)}(95)$ and $\mathrm{TAPS}_j^{(e)}(99)$ across j. We construct analogous quantities

$$M_j^{(a)}, \quad T_j^{(a,90)}, \quad T_j^{(a,95)}, \quad T_j^{(a,99)}$$

for agriculture. In all cases, higher values indicate a PU that is better ranked relative to other PUs on that dimension.

Step 5: Composite environmental and agricultural scores

Finally, we combine the percentile–ranked components into single composite environmental and agricultural scores for each PU. For environmental value, we define:

$$S_j^{(e)} = 0.40\, M_j^{(e)} + 0.25\, T_j^{(e,90)} + 0.20\, T_j^{(e,95)} + 0.15\, \big(T_j^{(e,99)}\big)^{1.5}, \tag{11}$$

where the weights put slightly more emphasis on average environmental quality and the broad high tail (90–95th percentile), while the exponent on $T_j^{(e,99)}$ gives additional, but non–dominant, credit to PUs containing exceptional extreme–tail areas.

We define the composite agricultural score analogously:

$$S_j^{(a)} = 0.40\, M_j^{(a)} + 0.25\, T_j^{(a,90)} + 0.20\, T_j^{(a,95)} + 0.15\, \big(T_j^{(a,99)}\big)^{1.5}. \tag{12}$$

By construction, $S_j^{(e)}, S_j^{(a)} \in [0,1]$, and they summarise, respectively, each PU's conservation attractiveness and agricultural opportunity cost. These composite scores are the quantities used as "environmental value" and "price" in the static knapsack problem and in the dynamic adversarial games.

A10 Replication Code and Supplementary Interactive Game

In addition to the robustness exercises presented above, we provide open replication materials to support transparency, reproducibility, and teaching applications.

First, we provide a standalone Python simulator, archived with a permanent DOI, together with a detailed User Manual:

- GitHub repository:
  https://github.com/dmweinhold/Conservation_Strategy_Simulation
- Zenodo archive: https://doi.org/10.5281/zenodo.17114490

The simulator can be used to reproduce all results reported in the paper and allows users to modify parameters such as the number of replications, correlation structure, allocation rules, and leakage rates.

Second, an interactive browser version of the Conservation Strategy Game, where users can play as either team against fixed strategies played by the computer, is available at:
https://dmweinhold.github.io/Conservation-Strategy-Game-Page/

Python Simulator & User Manual
https://github.com/dmweinhold/
Conservation_Strategy_Simulation

Interactive Browser Game
https://dmweinhold.github.io/
Conservation-Strategy-Game-Page/

Figure A10.1: QR codes linking to replication materials. The left panel links to the full simulator and documentation; the right panel links to the interactive browser game.