

Adversarial Procurement in Two-Value Space: Insights and Evidence for Conservation Siting*

Diana Weinhold[†] Lykke E. Andersen[‡]

[Click here for the latest version of this draft](#)

April 24, 2026

Abstract

We study a new general structural class we call *adversarial procurement in two-value space* where a mission-driven, budget-constrained buyer faces an active rival whose valuation of an asset sets its acquisition price. Strikingly, both threat-targeting and cost-effectiveness heuristics fail the focal agent in this setting. While prior literatures have established both failures under context-specific mechanisms — market-clearing feedback, stochastic threat, opportunity-cost heterogeneity — we identify rivalry plus two-value space as structural conditions under which the cost-effectiveness failure becomes generic in the benchmark environment. Conservation siting provides the most transparent instance, though the same structure arises across numerous domains, including public land acquisition, government procurement, and health systems.

The analysis proceeds computationally, formally, and on a real landscape; three nested lessons emerge. First, *rivalry changes the meaning of threat*: when rivals can re-target after being blocked, threat-targeting becomes fragile, and constraining re-targeting can paradoxically worsen outcomes by channelling uncontested assets into the focal agent’s portfolio without commensurate value. Second, *two-value space changes the meaning of cost*: when acquisition prices track the rival’s valuation, cheapness signals an uncontested asset, and a value-first rule asymptotically dominates the textbook cost-effectiveness ratio — a pattern we call the *knapsack reversal*. Its mechanism is *self-hedging*: the rival’s removal of the most expensive items absorbs the cost pressure that value-first’s selection would otherwise create. Finally, *valuation heterogeneity determines magnitude*: both distortions strengthen when the focal agent’s valuations vary meaningfully across items, and positive but imperfect correlation between the two value dimensions — the empirically common case — amplifies rather than attenuates them.

Formal theory establishes this as an asymptotic dominance theorem within the class (Theorems 4.1 and 4.2). A case study on Bolivia ($\rho \approx 0.29$, 390 planning units) confirms the qualitative rankings on a realistic landscape while revealing a 15–20% “disappointment gap” between static Marxan-style planning benchmarks and realised adversarial outcomes.

Key words: Adversarial procurement, Knapsack reversal, Conservation siting, Leakage

JEL: C72, C73, D61, Q57, Q24

*We thank Ben Balmford, Todd Kaplan, Lewis Grant, Timo Goeschl, Ben Groom, and Charlie Palmer for useful comments, Fabiana Argandoña for excellent research assistance, and the Natural Environment Research Council (NERC) for funding. This paper was developed with substantial assistance from AI systems: ChatGPT Pro (OpenAI) contributed to coding, proof & mathematical development, and auditing; Claude Opus (Anthropic) contributed to drafting, mathematical development, and proof repair; Gemini Pro (Google) contributed hostile auditing. The authors take sole responsibility for all claims.

[†]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

More than 190 countries have committed to protecting 30% of the Earth’s land and ocean by 2030, demanding an unprecedented acceleration of conservation efforts. But where should these protections go? Standard modelling frameworks treat this as a social planner’s optimisation problem, but conservation siting is often conducted within a problem structure we call *adversarial procurement in two-value space*: allocation is adversarial and sequential, with each side’s choices reshaping the other’s opportunity set; a blocked rival may re-target, making threat endogenous; and the focal agent’s cost is determined by the rival’s valuation rather than its own, so that cost signals not only expense but removal risk. This paper develops a framework for studying this structure, with conservation siting as the primary application, and shows that standard procurement heuristics fail under adversarial dynamics. Targeting the assets the rival most values (in conservation, prioritising sites under development threat) does not reliably outperform a value-first rule once the rival can adapt; constraining the rival’s ability to re-target (restraining *leakage*) can paradoxically increase welfare losses; and a standard benefit-cost ratio rule fails to maximize value for a fixed budget, a pattern we call the *knapsack reversal*. When price reflects the rival’s valuation, the “best deals” are the assets the rival does not contest, and apparent savings become a form of unilateral retreat. High price is not only a cost to Green but also the rival’s valuation — and, in a budgeted game, the rival’s cost of denial. Value-first forces contested ecological assets either to be protected or to be denied at rival expense; ratio-greedy may forgo them without making the rival pay.

Conservation provides the most empirically documented instance of adversarial procurement in two-value space. Land allocation is frequently adversarial: cooperative frameworks are difficult to sustain in practice, as dual-objective projects that simultaneously promote human development and conservation frequently achieve suboptimal environmental outcomes (Delacote et al., 2014; Amin et al., 2019), and political conflict around conservation is common. From Brazil, where conflicts between agribusiness interests and environmental agencies have produced abrupt swings in Amazon land-use policy (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 decisions often emerge from political arenas where actors with fundamentally different priorities struggle over institutional levers of power.

Conservation also exhibits the two-value-space structure transparently. Land prices track the developer’s agricultural valuation rather than the conservationist’s ecological payoff; the developer’s valuation drives both the price the conservationist pays and the probability that the site will be developed if the conservationist waits. Re-targeting by developers is empirically significant: estimates of leakage range from roughly 10–15% for Brazilian reforestation (Silva and Nunes, 2025) to 20–50% globally (Searchinger et al., 2018), with the degree shaped by enforcement, property rights, and market structure. Conservation is also unusual among adversarial procurement settings in that the counterfactual metric, called *additionality*, has become a binding evaluation criterion: carbon markets, payments for ecosystem services, and impact-evaluation methodologies increasingly channel funding toward sites where measured additionality is highest (Andam et al., 2008; Joppa and Pfaff, 2011; Wunder, 2015; Delacote et al., 2024), precisely the contested frontier where adversarial dynamics are most intense.

While conservation provides a particularly transparent instance of this structure, it is not unique: the same features arise wherever a mission-driven buyer deploys a budget against an active rival whose valuation helps set prices. Health systems value interventions for their health impact but compete for scarce clinical inputs against the private sector (Weinstein and Stason, 1977; Gold et al., 1996; Neumann et al., 2016); in government procurement and capital budgeting, projects are valued for their social benefit but compete against commercial returns (Brealey et al., 2020); and public land banks or housing authorities face prices driven by private development potential rather than by social value (Johnson and Hurter, 2000; Baum-Snow and Marion, 2009; Shavell, 2010; Miceli and Segerson, 2012).

Social planning approaches to these problems (e.g. Ando et al., 1998; Delacote et al., 2024; Gold et al., 1996; Brealey et al., 2020; Shavell, 2010) yield valuable insights into cost-effectiveness and institutional design, but assume away the adversarial dynamics that characterise many real-world procurement problems. Cost-effectiveness ratios, the workhorse allocation rule in all of these domains, reflect the fractional-knapsack optimality result (Dantzig, 1957): when a single agent fills a budget from a fixed menu, the greedy ratio rule is optimal. Our theorems show that this optimality is fragile under adversarial dynamics. When a rival’s purchases alter the menu, the ratio rule can be strictly dominated by a simpler rule that ignores costs entirely.

To study this class of problems, we develop a three-pronged approach. A computational

framework explores the high-dimensional interaction of strategy, rivalry, and leakage in two-value space, revealing economic mechanisms that static analysis obscures. Formal theory then establishes that the central finding — the knapsack reversal — is a structural property of the data-generating process, not an artefact of the simulation design. Finally, a case study on Bolivia tests whether the qualitative patterns survive on a realistic landscape.

The computational framework places conservationists (“Greens”) and developers (“Farmers”) in sequential competition over land that each side values differently. The analysis proceeds in two layers. In *Claims World*, we abstract from prices and budgets and study a pure rivalry problem with endogenous displacement. In *Budget World*, we add the central feature of procurement: Greens must buy sites at prices determined by Farmers’ agricultural value.

The simulations deliver three main findings. First, in *Claims World*, threat-targeting is much less robust than static conservation logic suggests. Blocking the rival’s preferred sites can create indirect gains when leakage is incomplete, but those gains erode as re-targeting rises, while the ecological penalty from neglecting intrinsically valuable sites accumulates regardless. Second, reducing leakage can perversely increase welfare losses: when developers are prevented from re-targeting, conservation may end up acquiring high-agricultural-value land without commensurate ecological gain. Third, once prices are introduced in *Budget World*, a standard ratio rule that ranks sites by ecological value per unit cost can perform worse than a value-first rule that ignores cost. This is the *knapsack reversal*. Its mechanism is what we call *self-hedging*: the same agricultural value drives Green’s acquisition cost, Farmer’s demand, and — in a budgeted game — Farmer’s cost of denial. Cheapness is therefore not a safety margin but a signal that the rival is uninterested; value-first either secures contested ecological value or forces the rival to spend scarce capacity denying it, while ratio-greedy often forgoes the same sites on cost-effectiveness grounds.

The fragility of threat-weighting under rivalry, and the knapsack reversal under adversarial procurement, are the paper’s two most surprising results. Two toy examples on the same six-plot board make these mechanisms concrete. The first illustrates the threat-chasing failure: how a pure threat-targeting rule sacrifices intrinsic ecological value for displacement gains that evaporate under adversarial rivalry with realistic leakage. The second illustrates the knapsack reversal: how a ratio-greedy rule loses high-value sites to development by deferring them on

cost grounds.

Consider a conservation buyer (“Greens”) choosing among items (plots) indexed by i , where each item has an environmental value e_i to Greens and an agricultural/development value a_i to a rival (“Farmers”). Greens must pay a price that reflects development value, so we take the price of item i to be a_i . Under budget parity, each side receives half the total agricultural value: $B_G = B_F = \frac{1}{2} \sum a_i = 15$. The players move sequentially (alternating turns), with Farmers moving first, until budgets are exhausted. Farmers follow their BAU rule, purchasing the highest- a remaining plot they can afford at each step. There are six plots with

$$(e_i, a_i) \in \{(10, 5), (8, 3), (7, 9), (5, 7), (3, 4), (2, 2)\}.$$

Farmers’ BAU ordering (by a) begins with $(7, 9)$, then $(5, 7)$ or the next affordable item.

Threat-chasing failure. Consider first a pure threat-chasing Green heuristic that selects the highest- a remaining plot (i.e., blocking the Farmer’s next BAU target). In round 1, Farmers select $(7, 9)$ and Greens threat-chase by selecting $(5, 7)$. In round 2, Farmers select $(10, 5)$ — now affordable since the two costliest items have been removed — and Greens select $(3, 4)$. In round 3, Farmer can no longer afford any remaining plot; Greens select $(8, 3)$ but cannot afford $(2, 2)$ with their remaining budget. Thus threat-chasing yields total environmental value $e = 5 + 3 + 8 = 16$.

By contrast, under a value-first rule Greens select the highest- e remaining plot on each turn. In round 1, Farmers select $(7, 9)$ and Greens select $(10, 5)$. In round 2, Farmers select $(3, 4)$ — since $(5, 7)$ now exceeds Farmer’s remaining budget — and Greens select $(8, 3)$. In round 3, Farmers select $(2, 2)$ and Greens select $(5, 7)$. Every plot is purchased and value-first yields $e = 10 + 8 + 5 = 23$.

Measured additionality (the increment in Greens’ environmental value relative to the BAU benchmark) is therefore lower under pure threat-chasing than under value-first. The intuition is that when development pressure (high a) is an imperfect guide to ecological value (high e), spending early turns primarily to protect developer targets crowds out the acquisition of intrinsically high- e parcels — and the loss is magnified in adversarial play because delayed high- e parcels may be removed before the next opportunity to act. Here the threat-chaser uses its

first turn on a mediocre plot ($e = 5$) while the ecological star ($e = 10$) is taken by Farmers in the next round. This is a failure of PSE that no amount of DLE can reliably compensate for: displacement gains arise only when leakage is incomplete, but the ecological cost of missing high- e sites is incurred regardless.

Knapsack reversal. Now consider a static cost-effectiveness (ratio-greedy) rule that ranks by e_i/a_i on the same board. In round 1, Farmers select (7, 9), after which Greens select (8, 3) (ratio $8/3 \approx 2.67$, the highest available). In round 2, Farmers select (10, 5), after which Greens select (2, 2) (ratio 1.00). In round 3, Farmer can no longer afford any remaining plot; Greens select (3, 4) (ratio 0.75) but cannot afford (5, 7) with their remaining budget. Thus ratio-greedy yields total environmental value $e = 8 + 2 + 3 = 13$ — even worse than threat-chasing (16), and far below value-first (23).

Starting from the same resources, ratio-greedy is ultimately *less* cost-effective than value-first because it uses early turns on high- e/a but low- e plots and thereby fails to pre-empt high- e plots that are removed under Farmers' a -based choices. The critical moment occurs in round 2: Farmers take the $e = 10$ plot — the best ecology on the board — which ratio-greedy deferred because its ratio ($10/5 = 2.0$) was below the top ratio ($8/3 \approx 2.67$). Value-first secured this plot immediately; ratio-greedy lost it to development. Moreover, Farmers capture substantially more ecological value when facing ratio-greedy ($e_{\text{Farmer}} = 7 + 10 = 17$) than when facing value-first ($e_{\text{Farmer}} = 7 + 3 + 2 = 12$): the ratio rule inadvertently feeds the rival better ecology. This is a failure of PSE: by deferring intrinsically high- e plots on cost-effectiveness grounds, the ratio-greedy rule loses them to development.

Formal theory establishes that the knapsack reversal is not a simulation artefact within the paper's calibrated Budget World theorem object. Under the paper's baseline data-generating process, Theorem 4.1 proves that, asymptotically, Max Environment strictly dominates Max Efficiency on Purchased Conservation (the total ecological value of sites the buyer secures) in the independent case, and Theorem 4.2 extends the dominance to $\rho = 0.3$, the positive ecological-agricultural correlation estimated on the Bolivian landscape. Both results hold under full leakage, so the advantage is not a residual-crediting or DLE result: value-first wins on the value of the sites it directly purchases. The rival remains budgeted, however, so the theorem still concerns a budgeted race over contested assets rather than a static menu. Constructive

counterexamples confirm that blanket dominance does not hold for every finite grid. These counterexamples rely on tight-budget mechanisms that become nonrepresentative under the benchmark sampling scheme as N grows.

The proofs reveal a clear mechanism. In an early selection phase, Max Environment secures the ecological elite while the ratio rule retreats into a cheap but ecologically mediocre region of the value space. In the post-prefix phase, the rival’s removals do not erase that lead; after the rival is exhausted, the ratio rule partially recovers and continues purchasing slightly longer, but not enough to reverse the ranking. The companion intuition guide (Weinhold, 2026) develops the economic content of these proofs through three structural features of self-hedging: a four-level decomposition, a cross-damage symmetry governing the baseline case, and a two-force balance between *selection amplification* and *cross-damage amplification* that shapes the correlated case. The theory also yields two scope insights. Positive correlation between ecological and agricultural value strengthens the reversal in the empirically relevant range rather than weakening it, while ecological compression weakens the reversal by shrinking the quality gap between top-ecological and top-ratio portfolios. This leads to a simple diagnostic, the *ecological selection premium* (ESP): the extent to which the ecologically best sites outperform the ratio-best sites in ecological value.

A case study on Bolivia tests these mechanisms on a realistic landscape. Using high-resolution biophysical data, we build a Bolivian planning board with 390 spatially heterogeneous units and estimate a positive ecological–agricultural correlation of approximately $\rho \approx 0.29$. The same qualitative ranking survives: value-first procurement continues to outperform both threat-targeting and static cost-effectiveness, and a Marxan-style static planning benchmark overstates realised conservation returns by roughly 15–20%.

These findings speak to a long-running practical and academic debate over whether scarce conservation resources should prioritise biodiversity, threat, or cost-effectiveness. Early reserve-selection work emphasised biodiversity value while largely abstracting from threat and cost heterogeneity (Margules et al., 1988; Faith, 1992). Later work advocated protecting high-value, high-risk “hot spots” (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). A separate tradition argued that cheaper and more remote sites yield greater cost-effectiveness (Ando et al., 1998; Naidoo et al.,

2006; Armsworth et al., 2006), while more recent work has again moved toward biodiversity-first approaches (Dinerstein et al., 2017, 2020, 2024). Meanwhile, the rise of additionality as a central evaluation criterion in carbon markets, payments for ecosystem services, and impact evaluation has channelled finance toward sites where measured counterfactual impact is highest (Engel et al., 2008; Aspelund and Russo, 2023)—precisely the contested settings where adversarial dynamics are most likely to matter. A small non-cooperative literature in conservation (Angelsen, 2001; Colyvan et al., 2011) underscores that conflict can drive outcomes far from the social optimum, but does not systematically compare the siting heuristics used in practice. Recent work on Markov-perfect equilibria in differential games (Jaakkola and Wagener, 2025), while exploring a setting that differs from ours in key respects (symmetric players, a single shared value dimension, continuous flow control rather than sequential discrete acquisition with displacement), reinforces a broader point: frameworks adopted for tractability can obscure economically important features of strategic interaction.

More broadly, the paper contributes a framework that is related to, but distinct from, several theoretical and applied literatures. Bilevel knapsack interdiction (Pferschy et al., 2021; Fischetti et al., 2019) and adversarial knapsack problems (Thakoor et al., 2025) study sequential adversarial budgeted acquisition within a single shared value space. Repeated second-price auctions with private budgets (Gummadi et al., 2012; Amar and Renegar, 2023) study rival-set prices but against a stochastic bid process rather than an active budget-constrained agent with displacement. In applied conservation economics, ratio-greedy failures have been established under several context-specific mechanisms: market-clearing feedback (Armsworth et al., 2006), output-price slippage (Wu et al., 2001), stochastic development threat (Costello and Polasky, 2004), and opportunity-cost heterogeneity (Ando et al., 1998; Naidoo et al., 2006). Our setting differs from each of these by combining rivalry with displacement, two distinct valuations per item, and a price mechanism that ties focal cost mechanically to the rival’s own valuation. The contribution of this paper is to isolate these three features as a structural source of the knapsack reversal, rather than treating ratio-greedy failure as a collection of unrelated contextual anomalies — with self-hedging as the organising mechanism producing it. Prior observations of ratio-greedy failure across these literatures become manifestations of a single structural pattern. To our knowledge, adversarial procurement in two-value space has not previously been studied as a named structural class.

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 decomposition. Section 3 reports the simulation results across leakage levels and developer behaviours. Section 4 extends the contest to Budget World, introduces the self-hedging mechanism and the knapsack reversal, states the paper’s asymptotic dominance theorems, explains the economic mechanism the proofs reveal, and characterises the ecological conditions under which the reversal operates. Section 5 turns to Bolivia: Section 5.1 examines historical patterns in land allocation, while Section 5.2 builds the Bolivian planning board, runs the strategy competitions, and compares dynamic portfolios to a static cost-effectiveness benchmark. Section 6 discusses policy implications, the role of correlation, and broader lessons for adversarial procurement. Robustness exercises appear in Appendices A1–A6; Appendix A8 contains the formal theory, including the asymptotic dominance theorems, counterexamples, and small-grid benchmarks; and supplementary material on Bolivia appears in Appendix A9. In a short companion brief to this paper, Weinhold (2026) provides a guide to the intuition behind the theoretical proofs presented in Appendix A8. Replication materials and code are available via GitHub and archived on Zenodo (DOI: 10.5281/zenodo.19613421)

2 Simulation Framework: Baseline Claims World

We now formalize the adversarial contest previewed in the introduction as a Monte Carlo simulation framework in which two teams — Greens (conservationists) and Farmers (developers) — compete over a grid of land plots. The computational framework allows us to isolate two distinct mechanisms within *adversarial procurement in two-value space*: in Claims World, developed in this section, we isolate the role of adversarial rivalry with displacement, abstracting away from the two-value price mechanism. In Budget World, developed in Section 4, we restore the full structure to reveal a second mechanism, self-hedging, which emerges from the combination of adversarial rivalry and two-value space. Together, the Claims World and Budget World layers let us separate what rivalry does on its own from what the two-value price mechanism adds.

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 plot is pre-emptively claimed by the Greens. Each Farmer strategy determines a fixed “business-as-usual” (BAU) set at the start of the game — for a naïve profit maximiser, this is simply the list of plots ranked from highest to lowest agricultural value a_i , targeted in that order. BAU is a static concept: it describes what the Farmer *would* do in the absence of Green intervention, and it does not change as the game unfolds. Leakage governs what happens when Green disrupts this plan. 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 formally in Section 2.2.4), and conservation gains from the residual left-over plots (what we call the *Displacement-Leakage* effect, also defined in Section 2.2.4), 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 Claims World game is an alternating draft over indivisible plots. Teams take turns claiming currently available plots until their claims are exhausted. We begin with a 10×10 grid of plots, each assigned an environmental value and an agricultural value.

Let

$$P = \{p_i\}_{i=1}^N, \quad N = 100,$$

denote the set of plots. Each plot p_i has two attributes:

$$e_i \geq 0 \quad (\text{environmental value}), \quad a_i \geq 0 \quad (\text{agricultural value}).$$

The pair (e_i, a_i) is generated from a Gaussian copula so that the dependence between the two value dimensions can be varied by the user. Specifically,

$$(z_{e,i}, z_{a,i}) \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}\right),$$

where ρ is the latent Gaussian-copula correlation parameter. These latent normal draws are

mapped to uniforms by

$$u_{e,i} = \Phi(z_{e,i}), \quad u_{a,i} = \Phi(z_{a,i}),$$

and then linearly rescaled to $[0.1, 10.0]$:

$$e_i = 0.1 + 9.9 u_{e,i}, \quad a_i = 0.1 + 9.9 u_{a,i}.$$

This construction ensures that the marginal distributions of e_i and a_i are uniform on $[0.1, 10.0]$, with dependence induced through the Gaussian copula. 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 τ is

$$\tau = \frac{2}{\pi} \arcsin(\rho).$$

Appendix [A5](#) shows that our results are robust to heavy-tailed environmental draws.

The grid structure is purely practical: there is no spatial information in the baseline Claims World grid. The framework can map to any geographical arrangement or size distribution of plots. We interpret e_i and a_i as the net present values of protection and development, respectively, including any dynamic externalities. More generally, we assume that the rank ordering of these values summarizes the multidimensional considerations that drive each team's siting preferences.

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

Let T_f and T_g denote the Farmer and Green teams. Each team $T \in \{T_f, T_g\}$ receives an initial endowment of integer claims C_T , with

$$C_f + C_g = N.$$

By varying the allocation of claims between Greens and Farmers, we simulate different degrees of political power.

The game runs for R rounds. In each round r , Farmers move first, followed by Greens. The convention that Farmers move first is the more conservative test of the value-greedy Green strategy *Max Environment*; tiny-grid tests (Appendix A8.9) illustrate that allowing Greens to move first gives it a slight initial advantage, though this dissipates rapidly as grid size grows.

For each team T , let $C_T^r \in \mathbb{N}_0$ denote the integer number of claims available in round r , with

$$\sum_{r=1}^R C_T^r = C_T.$$

In the equal-power baseline with C_T divisible by R , we have

$$C_T^r = \frac{C_T}{R} \quad \text{for each } r = 1, \dots, R.$$

When C_T/R is non-integer (as in the unequal-power exercises of Appendix A4), each team claims $\lfloor C_T/R \rfloor$ plots per round, with the remaining

$$C_T - R\lfloor C_T/R \rfloor$$

extra claims distributed one per round across the earliest rounds until exhausted. Thus C_T/R describes the average claiming rate, while the simulation implements the exact integer schedule.

Let $S_f^r \subseteq P$ and $S_g^r \subseteq P$ denote the sets of plots claimed by Farmers and Greens in round r . The cumulative claimed sets after round r are

$$S_f^{\leq r} = \bigcup_{k=1}^r S_f^k, \quad S_g^{\leq r} = \bigcup_{k=1}^r S_g^k.$$

A plot can be claimed at most once, so

$$S_f^{\leq r} \cap S_g^{\leq r} = \emptyset \quad \forall r.$$

Let

$$U^{r,0} := P \setminus (S_f^{\leq r-1} \cup S_g^{\leq r-1})$$

denote the set of plots available at the start of round r .

When a team has more than one claim in a round, it selects plots sequentially from the current within-round availability. We therefore write $s_{T,\ell}^r$ for team T 's ℓ -th claim in round r .

The within-round available sets are defined recursively by

$$U_f^{r,0} := U^{r,0}, \quad U_f^{r,\ell} := U_f^{r,\ell-1} \setminus \{s_{f,\ell}^r\}, \quad \ell = 1, \dots, C_f^r,$$

and

$$U_g^{r,0} := U_f^{r,C_f^r}, \quad U_g^{r,\ell} := U_g^{r,\ell-1} \setminus \{s_{g,\ell}^r\}, \quad \ell = 1, \dots, C_g^r.$$

Thus the next round begins with

$$U^{r+1,0} = U_g^{r,C_g^r}.$$

The round- r claim sets are therefore

$$S_f^r = \{s_{f,1}^r, \dots, s_{f,C_f^r}^r\}, \quad S_g^r = \{s_{g,1}^r, \dots, s_{g,C_g^r}^r\}.$$

The strategy rules below define the *next within-round claim*. When $C_T^r > 1$, the same rule is applied recursively to the updated within-round available set until the round quota is exhausted.

In the baseline theorem object of Appendix A8, $C_f^r = C_g^r = 1$, so this reduces to the alternating one-pick game.

2.2.1 Leakage

Let $\sigma_f^{\text{BAU}} = (q_{f,1}, \dots, q_{f,C_f^r})$ denote the Farmer *business-as-usual* (BAU) sequence: the ordered list of plots that the Farmer's strategy would claim if it could use all its claims without any interference from the Greens. This sequence is determined entirely by the Farmer's strategy and the initial grid, and is fixed at the start of play. We write

$$S_f^{\text{BAU}} := \{q_{f,1}, \dots, q_{f,C_f^r}\}$$

for the corresponding BAU set.

During the game, if the Greens claim a plot in S_f^{BAU} , the Farmer must deviate from this

plan. If conservation efforts face 100% leakage, the Farmer can simply claim an alternative plot elsewhere in the grid. Thus under full leakage both teams spend all their claims, and all plots in P are allocated by the end of the game.

We then vary the degree of leakage, denoted

$$\text{Leakage} \in [0, 1],$$

by reducing the Farmer's remaining unspent claims proportionally for every plot $p_i \in S_f^{\text{BAU}}$ claimed by Greens. For example:

- if Leakage = 0.5, Farmers lose 1 unspent claim for every 2 BAU plots pre-empted by Greens;
- if Leakage = 0, Farmers lose 1 unspent claim for every 1 BAU plot pre-empted by Greens.

The game concludes when both teams have exhausted their claims, or when no further plots are eligible for claim. If Leakage < 1 and Greens have pre-empted Farmer BAU plots, there may be residual unclaimed plots at the end of the game. We define the final residual set by

$$U^{\text{final}} := P \setminus (S_f^{\leq R} \cup S_g^{\leq R}).$$

These plots are credited to Greens in the final-round accounting convention described below.

2.2.2 Farmer Strategies

Farmers are static profit maximisers who choose plots to maximise their own agricultural value. They differ in whether they model Green's behaviour when choosing:

1. **Naïve Profit Maximiser (rival-blind).** The Farmer does not observe the ecological value of plots and does not treat the conservation agency as a rival. At each opportunity, Farmer simply claims the plot with the highest remaining agricultural value:

$$s_{f,\ell}^r \in \arg \max_{p_i \in U_f^{r,\ell-1}} a_i.$$

“Naïve” thus refers to this informational and strategic limitation, not to any departure from canonical profit maximisation.

2. **Strategic Profit Maximiser (rival-aware).** The Farmer recognises that Green may target plots with high ecological value and adjusts accordingly, trading off own agricultural value against the risk that a preferred plot is pre-empted by Green. Strategic Farmers may therefore choose plots with slightly lower agricultural value but higher ecological value in order to secure ecologically attractive plots before Green can claim them.

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 as the risky set \mathcal{R}_τ . In each claim they choose the available plot with highest a_i inside \mathcal{R}_τ ; once \mathcal{R}_τ is exhausted they continue by descending a_i in the safe set \mathcal{R}_τ^c :

$$s_{f,\ell}^r \in \begin{cases} \arg \max_{p_i \in U_f^{r,\ell-1} \cap \mathcal{R}_\tau} a_i, & \text{if } U_f^{r,\ell-1} \cap \mathcal{R}_\tau \neq \emptyset, \\ \arg \max_{p_i \in U_f^{r,\ell-1} \cap \mathcal{R}_\tau^c} a_i, & \text{otherwise.} \end{cases}$$

This stylized behavior has clear empirical parallels in settings where landowners anticipate conservation interventions and adjust accordingly. For example, studies of the U.S. Endangered Species Act document pre-emptive timber harvests or habitat destruction to avoid regulatory restrictions (Lueck and Michael, 2003; Langpap and Wu, 2017).

2.2.3 Conservation Strategies

Within round r , Greens choose sequentially from the set remaining after the Farmers’ within-round claims. Let $U_g^{r,\ell-1}$ denote the available set just before the Green’s ℓ -th claim. The rules below define $s_{g,\ell}^r$, the next Green plot selected within round r .

1. **Max Environment (value-greedy strategy).** Greens claim the available plot with the highest environmental value:

$$s_{g,\ell}^r \in \arg \max_{p_i \in U_g^{r,\ell-1}} e_i.$$

2. **Block Farmers (threat-greedy strategy).** Greens claim the available plot with the

highest agricultural value:

$$s_{g,\ell}^r \in \arg \max_{p_i \in U_g^{r,\ell-1}} a_i.$$

3. **Hot Spot (product-greedy strategy).** Greens claim the available plot maximizing the product of environmental and agricultural value:

$$s_{g,\ell}^r \in \arg \max_{p_i \in U_g^{r,\ell-1}} (e_i a_i).$$

4. **Maximize Difference (net social benefit strategy).** Greens claim the available plot maximizing the difference between environmental and agricultural value:

$$s_{g,\ell}^r \in \arg \max_{p_i \in U_g^{r,\ell-1}} (e_i - a_i).$$

5. **Random Protection.** Greens claim a plot uniformly at random from the currently available set:

$$s_{g,\ell}^r \sim \text{Unif}(U_g^{r,\ell-1}).$$

2.2.4 Baseline Outcomes

We conduct comparative-statics exercises 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.

For terminal allocations we abbreviate

$$S_f := S_f^{\leq R}, \quad S_g := S_g^{\leq R}.$$

1. **Final Conservation Score.**

The final conservation score is the sum of the environmental values in the plots claimed by

Greens during the game, plus any residual unclaimed plots left undeveloped at the end:

$$C_{\text{final}} = \sum_{p_i \in S_g^{\leq R}} e_i + \sum_{p_i \in U^{\text{final}}} e_i.$$

Let $S_f^{\text{BAU}} \subseteq P$ denote the Farmer BAU set defined on the initial grid. Given this BAU pattern, let $S_g^{\text{BAU}} \subseteq P$ be the set of plots 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 portfolio 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 through effective targeting of Farmer BAU plots under incomplete leakage:

$$DLE = \sum_{p_i \in U^{\text{final}}} e_i.$$

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 Greens minus the conservation score they would have obtained had Farmers been allowed to claim all their BAU plots:

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

$$\text{Additionality} = C_{\text{final}} - C_{\text{BAU}}.$$

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:

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

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

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 compare broad classes of procurement heuristics rather than to characterise the equilibrium of the full dynamic game. This is a deliberate methodological choice. The exact discrete dynamic program depends on the full remaining set of plots, remaining budgets or claims, and the history-dependent BAU/burn state, so the state space grows combinatorially with grid size; but even if the equilibrium were tractable, it would answer a different question from the one we pose. The heuristics we study span a strategy continuum from pure threat-targeting (blocking the rival's preferred assets) through balanced approaches to pure value-maximisation (securing the highest-payoff assets regardless of the rival), with cost-effectiveness entering in Budget World as a budget-constrained variant of value-maximisation. The advantage of a heuristic horserace is that it reveals the deep principles governing this continuum: how the balance between intrinsic value, threat, and cost-awareness shifts the focal agent's outcomes as conditions change.

That said, formal theory plays an important complementary role. Appendix [A8](#) gives constructive counterexamples establishing that no blanket finite- N dominance theorem holds be-

tween the heuristics, tiny-grid backward-induction benchmarks that verify heuristic performance against equilibrium play, and — most importantly — two asymptotic dominance theorems proving that the knapsack reversal is a structural property of the baseline data-generating process, not an artefact of particular grid sizes or parametrisations. The theory thus confirms that the simulation rankings reflect a genuine economic mechanism rather than an artefact of the simulation 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 $\text{corr}(e, a) = 0$. 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).

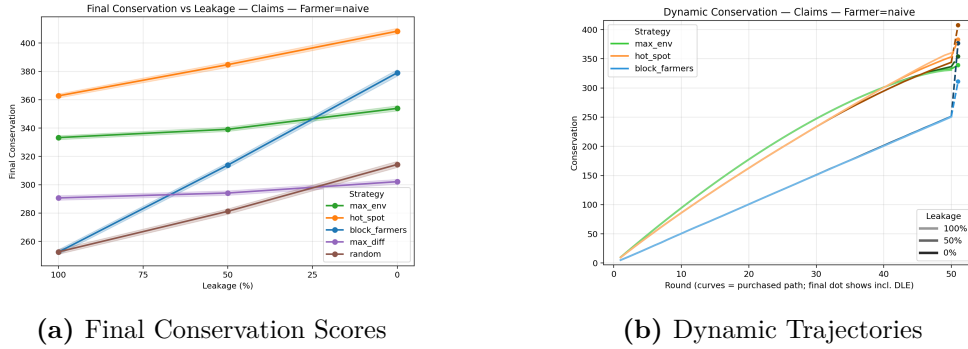
3.1 Conservation outcomes of Green strategies

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.

As leakage falls, Farmers lose the ability to replace BAU plots captured by Greens, and



(a) Final Conservation Scores

(b) Dynamic Trajectories

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

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 feasibil-

ity 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 $\text{corr}(e, a) = 0$.

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 .

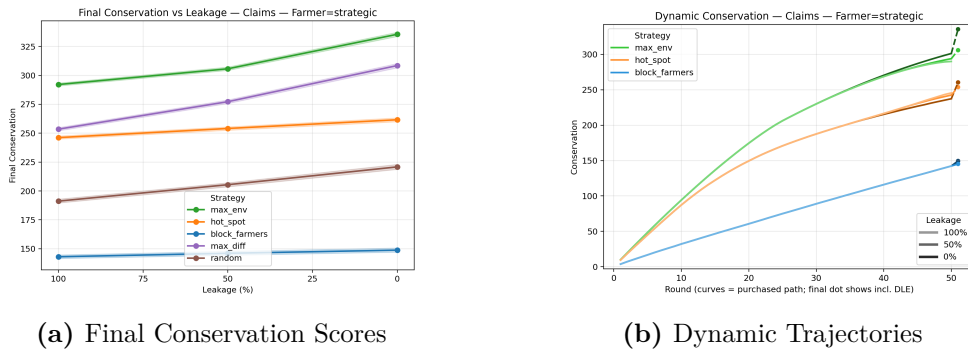


Figure 2: Claims World: Green Strategies and Strategic Farmers — Conservation Scores

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.

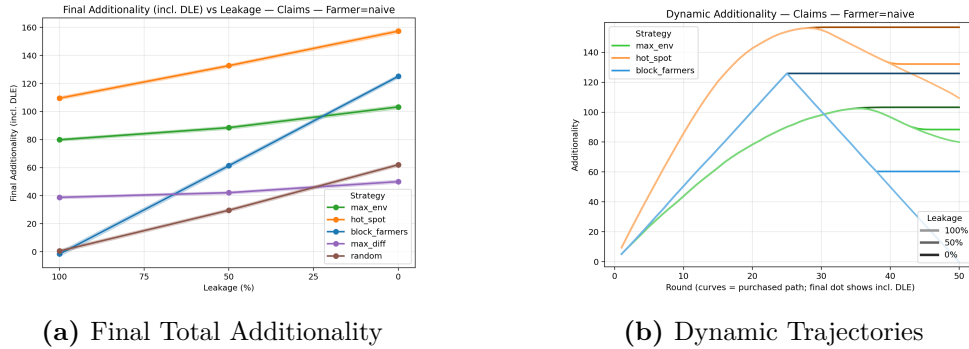


Figure 3: Claims World: Green Strategies and Naïve Farmers — Additionality

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.

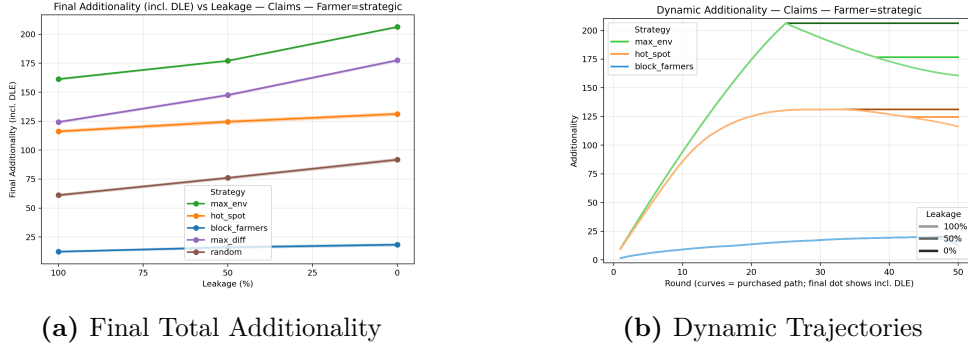


Figure 4: Claims World: Green Strategies and Strategic Farmers — Additionality

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.

The Strategic-Farmer result foreshadows the broader contestability logic that becomes central in Budget World. When Farmers are rival-aware, high-environmental plots are not merely valuable to Green — they are also targets that may be lost if Green delays. Claims World has no prices, so it cannot generate the knapsack reversal itself, but it already exposes the dynamic mistake behind threat-weighted rules: chasing the rival’s BAU set gives a rival-aware opponent time to pre-empt the ecological frontier. Budget World then converts this contestability problem into a price-signal problem.

Synthesis: Claims World Conservation and Additionality

Three patterns emerge from the Claims World results, each illuminated by the PSE/DLE decomposition. First, despite appearing *ex ante* to maximize additionality, Block Farmers performs poorly: with uncorrelated e and a and iterative play, it diverts effort to low-environmental plots whose PSE contribution is weak. Under high leakage, Farmers simply re-target to next-best agricultural plots (sometimes with high e), eroding DLE without any compensating PSE gain; under strategic Farmer behaviour, developers explicitly chase high- e plots, further undermining measured additionality. Second, the mechanism behind this failure is endogenous re-targeting:

once a preferred plot is blocked, developers reallocate, so “threat” is not fixed but adapts to conservation actions—a dynamic the PSE/DLE lens makes visible but that aggregate outcome measures obscure. Third, while Hot Spot and Max Environment are the best performers overall, their relative advantage depends on developer behaviour through distinct channels: Max Environment dominates against Strategic Farmers because it sustains the highest PSE and, because the Strategic Farmer BAU set is loaded with high- e plots, also reaps a larger DLE; Hot Spot can edge ahead against Naïve Farmers primarily through DLE at contested frontiers, though its PSE advantage is modest.

Appendix A8.9 provides a useful robustness check against the specific Farmer heuristics used in the simulations. On the reported tiny-grid backward-induction exercises, Hot Spot often lies closest to the equilibrium Green payoff in Claims World, Max Environment is usually next best, and Block Farmers performs poorly. These calculations are suggestive rather than definitive for larger boards, but they point in the same qualitative direction as the Monte Carlo results. Note that as $\rho \rightarrow 1$ the strategy distinctions narrow: under perfect positive correlation, high- e plots are also high- a plots, so Max Environment, Block Farmers, and Hot Spot converge to the same selection rule; under perfect negative correlation, developer and conservation targets are fully separated, eliminating the adversarial tension that drives our main results. The intermediate correlations explored in Appendices A2 and A3 preserve the qualitative rankings, and the Bolivian case study ($\rho = 0.29$) sits comfortably in the interior.

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.

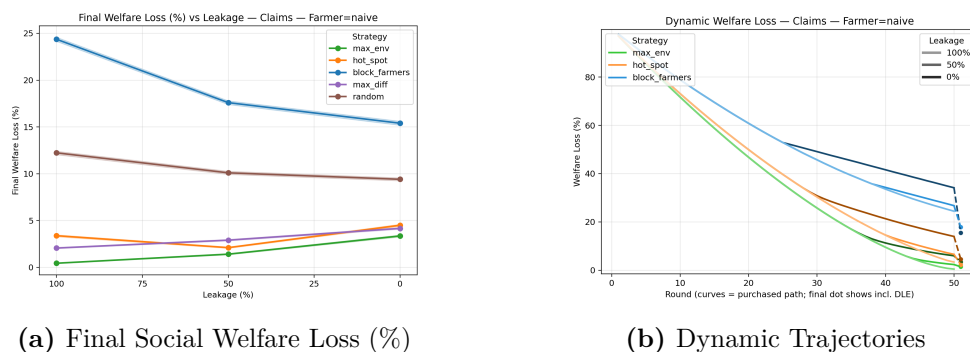


Figure 5: Claims World: Green Strategies and Naïve Farmers — Social Welfare Loss (%)

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.

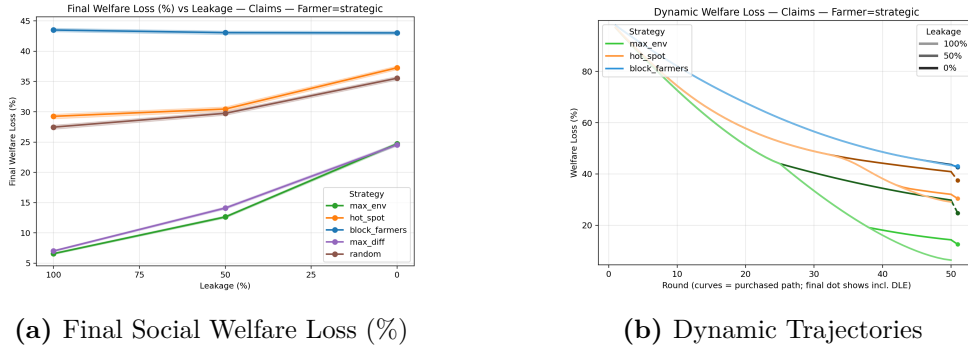


Figure 6: Claims World: Green Strategies and Strategic Farmers — Social Welfare Loss (%)

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

Claims World isolated the role of rivalry in a setting where both teams allocate discrete claims, abstracting from cost. We now extend the contest to *Budget World*, where both Farmers and Greens purchase plots subject to fixed initial budgets and plot-specific costs equal to agricultural values a_i . This preserves the adversarial contest while introducing the price mechanism that

drives the knapsack reversal previewed in the introduction.

The formal set-up is outlined in Appendix A7 and follows the logic of Claims World, with one key change: the price of each plot is set equal to the Farmers’ valuation, a_i . Because Greens pay a price determined by Farmers’ agricultural valuation rather than by their own ecological valuation, the two value dimensions play distinct roles in the contest: price tracks the rival’s valuation, while Greens’ own payoff depends on the ecological value e_i . We call this the *two-value-space* structure, and it distinguishes Budget World from standard one-space contest and interdiction/knapsack settings in which prices are set by the focal agent’s own valuation or by a single shared value. The structure is the source of the knapsack reversal: the ratio-greedy rule treats high prices as costs to be avoided, but in this two-space setting high prices also signal high removal risk, making avoidance a form of unilateral retreat. In our baseline Budget World analysis, environmental and agricultural values are drawn independently, so the price faced by Greens is fully informative about the Farmers’ payoff but carries no information about Greens’ own payoff; correlation between e and a changes the informational content but not the basic structure, and the knapsack reversal in fact strengthens as correlation rises into the empirically common moderate-positive range (Theorem 4.2; see Section 4.1.1). 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 adapt the Claims World logic to a budget-constrained setting. The key difference concerns the Strategic Farmer’s “risky” set: in Claims World, this is

defined by a simple rank cutoff (the top- C_g plots by environmental value), whereas in Budget World it is defined by the set of plots that a Green Max Environment rule could actually afford to purchase under B_g^0 . This affordability conditioning is the natural analogue: high- e plots whose agricultural value exceeds the Green budget are not credible Green targets and need not be pre-empted. In both Worlds, the Farmer purchases within the risky set by descending a_i , then fills from the safe set by descending a_i . Here and throughout the Budget World appendix,

$$\text{Residual} := P \setminus \left(S_f^{\leq R} \cup S_g^{\leq R} \right).$$

Formally:

1. **Naïve budgeter:** purchase the highest- a_i plot feasible within B_f^0 .
2. **Strategic budgeter:** identify “risky” plots as those that would be purchased by a Green Max Environment strategy operating under B_g^0 —that is, the set of affordable plots a greedy-by- e buyer would acquire given the Green budget. This differs from the Claims World definition, where the risky set is simply the top- $\lceil \tau N \rceil$ plots by environmental value regardless of cost; in Budget World, affordability filters out high- e plots whose price a_i would prevent Green acquisition, removing them from the set the Farmer needs to pre-empt. Purchase within the risky set in descending a_i until budget or set is exhausted, then continue with 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 show that the core Claims World findings carry through to the paper’s benchmark Budget World exercises. 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 performs best on long-run purchased

conservation outcomes. Likewise, strategic behavior by Farmers sharpens relative performance differences. A parallel set of tiny-grid equilibrium calculations (Appendix A8.9) points in the same direction: Max Environment looks weak on 2×2 but improves sharply by 3×3 , whereas Block Farmers performs badly throughout. We treat those calculations as supportive numerical evidence, not as a general proof.

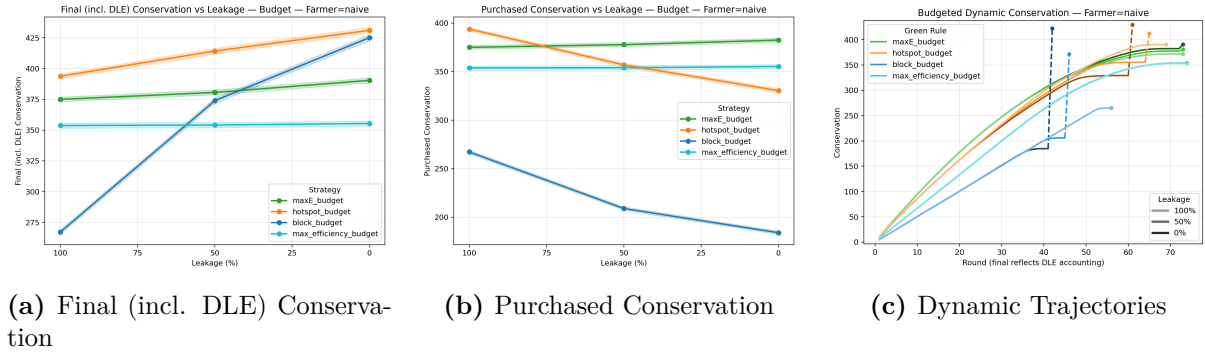


Figure 7: Budget World Green Strategies and Naïve Farmers: Conservation Scores

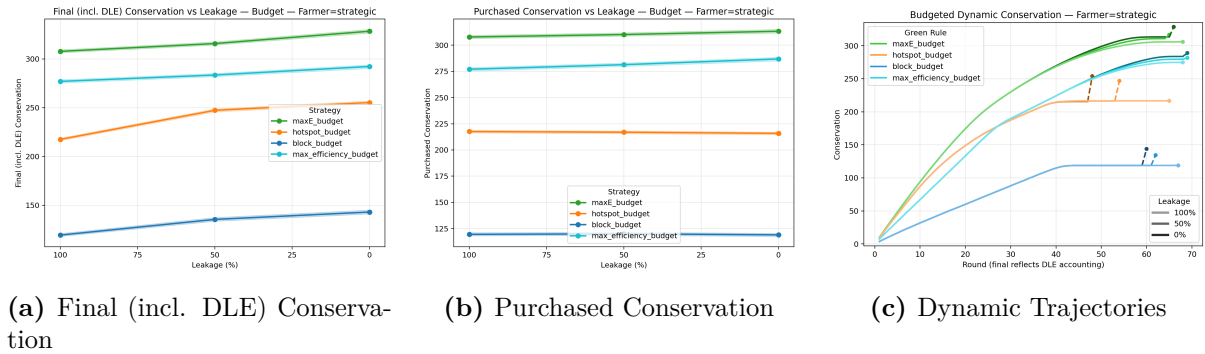


Figure 8: Budget World Green Strategies and Strategic Farmers: Conservation Scores

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. In the baseline simulations, the knapsack reversal previewed in the introduction is borne out quantitatively: across 500 replications at each leakage

level, Max Environment beats Max Efficiency in essentially 100% of cases, with a mean gap of roughly 20–30 units of e . The ranking is preserved at every reported leakage level, and the two strategies’ PC curves are approximately parallel across L , suggesting that the effect is driven mainly by selection rather than by residual crediting or extreme sensitivity to leakage.

The reversal is not confined to the Naïve Farmer benchmark. It also appears against Strategic Farmers, whose BAU set is explicitly loaded toward plots that a value-first Green buyer could credibly acquire, and it appears at both $\rho = 0$ and $\rho = 0.3$ in diagnostic runs. The distinction between the two Farmer types is informative. Against Naïve Farmers, the ME–MX gap is largest when leakage is low, consistent with a blocking channel: Green’s purchases of Farmer-BAU plots reduce future Farmer capacity, and this effect is strongest when leakage does not reallocate displaced development elsewhere. Against Strategic Farmers, the gap is substantially stable across leakage levels and largest under high leakage, consistent with a contestation channel: delaying high- e acquisitions is costly when the rival is also looking for them. The computational pattern is therefore not an artefact of myopic Farmer behaviour; different adversarial channels reinforce value-first play in different leakage regimes.¹

Budget World: Welfare Loss outcomes

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

4.1.1 The knapsack reversal: theory

This section establishes theoretically that the surprising finding reported above — that a simple value-maximisation rule reliably outperforms the textbook cost-effectiveness heuristic under adversarial dynamics, a pattern we call the *knapsack reversal* — is not an artefact of the simulation design. The two theorems below prove asymptotic dominance of Max Environment

¹In additional diagnostic runs, a costless policy-targeting adversary — one that deletes Green’s next intended purchase for free — largely eliminates the ME advantage. This is consistent with the self-hedging interpretation: the reversal survives budgeted strategic denial but not free targeted denial. We treat this diagnostic as a stress test rather than an economically meaningful Farmer strategy; a budgeted analogue would correspond to a rival whose objective is to harm Green rather than to acquire land, which falls outside the profit-maximising class of rivals the paper studies.

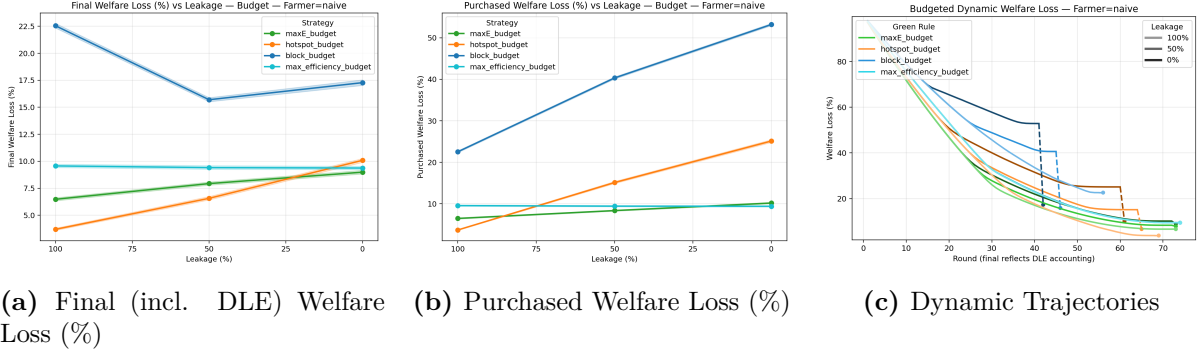


Figure 9: Budget World Green Strategies and Naïve Farmers: Social Welfare Loss (%)

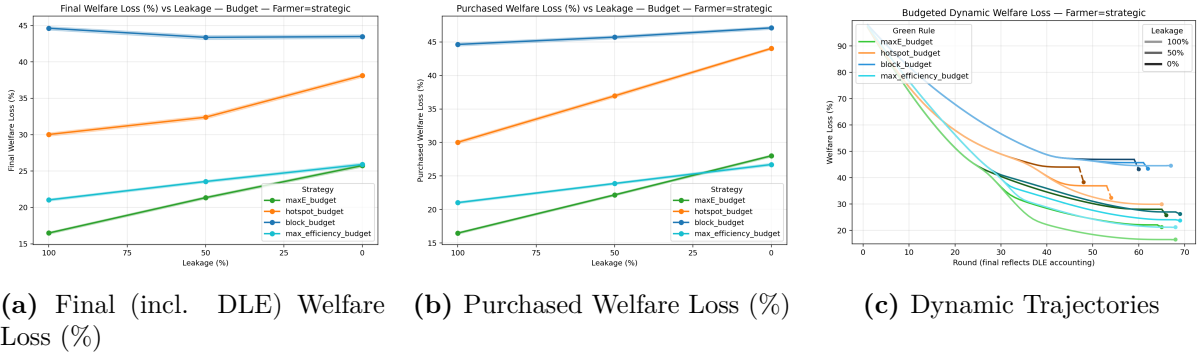


Figure 10: Budget World Green Strategies and Strategic Farmers: Social Welfare Loss (%)

over Max Efficiency against a naive max- A Farmer at two representative latent correlations: $\rho = 0$ in closed form, and $\rho = 0.3$ via a computer-assisted validated certificate. Simulations show the reversal persisting across a wider range of correlations and against the Strategic Farmer benchmark, but the formal guarantees are restricted to the two correlation points at which the fluid-limit machinery closes. Appendix A8.3 provides the proof roadmap and constant ledger; the full baseline proof is in Appendices A8.4–A8.5, and the $\rho = 0.3$ extension is in Appendix A8.6. We focus here on the economic content.

Both results are proved for Purchased Conservation under full leakage — the setting in which indirect displacement effects are absent by construction. The advantage is therefore not a residual-crediting or DLE result: value-first wins on the value of the sites it directly purchases. The rival remains budgeted, however, so the theorem still concerns a budgeted race over contested assets rather than a static menu.

The mechanism, in one sentence. Ratio-greedy fails because the selection criterion itself biases it into the part of the distribution the rival does not contest: the cheap-for-their-ecology

corridor of the landscape. Value-first competes head-to-head with the Farmer in the contested upper tail, where ecological value is highest. The rest of this subsection develops the two halves of this claim — why ratio-greedy selects a systematically weaker portfolio, and why value-first does not pay for competing directly — before turning to the conditions under which the mechanism operates.

Why ratio-greedy selects a weaker portfolio. The ratio e/a is high either because e is high or because a is low. Under independence, conditioning on a high ratio pulls strongly toward low- a items and only imperfectly toward the ecological frontier. The selected sites are therefore not necessarily low-value in absolute terms; indeed their average e may exceed the unconditional mean. The selection loss is relative to the top- e portfolio that value-first would have bought. Ratio-greedy pays for cheapness by moving away from the ecological elite and into a less contested corridor of the landscape.

Positive correlation amplifies this wedge.² When e and a move together, an item can achieve an unusually high e/a ratio only by having agricultural value that is low relative to its ecological value. These residual-cheapness sites are less attractive to the Farmer than the high- e , high- a frontier plots that value-first targets. The ratio portfolio can still be good in absolute ecological terms, but the gap to the ecological elite grows. The gap between the top- e and top- e/a portfolios (what we call the *ecological selection premium* below) increases from 0.219 at baseline to 0.309 at $\rho = 0.3$, and simulations indicate that the normalised gap is hump-shaped in ρ , peaking near $\rho \approx 0.7$ and remaining positive across the simulated range (Figure 11). The hump arises from two rising forces with opposing effects on the gap. *Selection amplification* strengthens the knapsack reversal: as ρ increases, the gap between ratio-greedy’s portfolio and the ecological elite widens because high-ratio items are increasingly those with anomalously low a relative to their e . *Cross-damage amplification* weakens it: Farmer’s max- A deletions become more ecologically damaging to value-first as ρ rises, because the most expensive items are increasingly also the most ecologically valuable. The two forces balance near $\rho \approx 0.7$; beyond that point, cross-damage amplification dominates. The policy implication is direct: cost-effectiveness mandates are most counterproductive precisely on the contested landscapes

²A natural prior could be that positive correlation should help ratio-greedy, because the highest-ecological sites are now also expensive and a “best bang for the buck” rule seems better suited to finding hidden gems. The theory shows the opposite.

where ecological and agricultural values are positively aligned — the empirically common case.




Figure 11: Normalised Purchased Conservation gap ($\Delta/N = [\text{PC}_N^{\text{ME}} - \text{PC}_N^{\text{MX}}]/N$) as a function of latent correlation ρ . The simulated gap is positive across the tested range and peaks near $\rho \approx 0.7$. Solid triangles mark the two correlation values at which asymptotic dominance is formally proved (Theorem 4.1 at $\rho = 0$; Theorem 4.2 via computer-assisted validated certificate at $\rho = 0.3$). The vertical dashed line indicates Bolivia’s estimated correlation ($\rho \approx 0.29$).

Why value-first does not pay for competing. The harder direction of the argument is showing that value-first’s head-to-head competition with the Farmer does not destroy its lead through higher continuation costs or faster budget exhaustion. Here a single economic idea — *self-hedging* — organises the answer across both correlations.³

At $\rho = 0$, self-hedging takes a particularly clean form. Because $E \perp A$, value-first’s picks are effectively random with respect to cost: Green removes the best remaining ecological site without any bias toward cheap or expensive items. The result is a perfect symmetry in the fluid limit — the Farmer’s frontier on the cost dimension and Green’s frontier on the ecological dimension deplete the pool at identical rates (Lemma A8.8). An immediate consequence is that Farmer’s per-turn ecological capture from value-first’s pool exactly equals Green’s per-turn agricultural cost from Farmer’s pool. The two agents are racing on symmetric but separate dimensions, and the race is fair. The asymptotic gap is then the difference between what MX cedes by design ($\mu_E \phi$: the ecological value of the upper- A fraction it declines to contest) and

³The self-hedging mechanism was hypothesised as the economic story behind the knapsack reversal early in this project, before the formal theorems were proved. The proofs subsequently confirmed it as a formal mathematical property of the game, not merely a post-hoc reading of simulation patterns. A four-level development is given in the companion intuition guide (Weinhold, 2026).

what ME loses by racing ($S_G(\tau)$: its cumulative agricultural cost through Farmer exhaustion). This identity is proved in Theorem 4.1⁴.

At $\rho = 0.3$, the symmetry breaks: Farmer’s max- A picks now correlate with high E , so Farmer captures more ecological value against ME than Green’s agricultural cost alone would predict. The correlated proof therefore cannot recycle the baseline symmetry identity. Instead it tracks the actual doubly truncated survivor pool in latent (X, Y) coordinates, proves a density-dependent fluid limit for Green’s max- Y path, and uses a stopping-time affordability bootstrap: the validated fluid budget remains bounded away from zero through the certified horizon, so an affordability failure would contradict uniform convergence of the budget coordinate.

The max- A structure of the Farmer is important for the proof, but not because the computational reversal appears only under myopic behaviour. Rather, the max- A rule exposes the self-hedging mechanism in a form that is mathematically tractable: Farmer interference has a predictable distributional structure, because the same deletion that removes a contested item also removes an expensive item, allowing the proof to control both value and affordability through the fluid-limit state. A fully strategic Farmer could target Green’s continuation value in ways that do not satisfy the same monotone cost lemma, so the theorem should not be read as a general strategic-Farmer result. The Budget World simulations, however, indicate that the reversal persists against the Strategic Farmer benchmark. The formal theorem therefore proves a clean sufficient mechanism within a wider computational pattern, rather than exhausting the mechanism’s empirical reach.

The combination of a large prefix quality advantage (Figure 11) and explicit tracking of Green’s actual per-turn ecological value through the fluid-limit ODE is what makes the correlated extension work. Theorem 4.2 formalises the result; the full proof architecture, which is genuinely more involved than the symmetric baseline, is given in Appendix A8.6.

Why the early lead survives. Figure 12 makes the logic visual. The top panel shows that the theorem is *not* a claim of pointwise dominance: the per-turn ecological value curves cross,

⁴The baseline theorem admits two additional proof architectures that yield more conservative lower bounds on the same asymptotic gap: a closed-form fluid-limit integral (certified margin ≥ 0.168) and a worst-case deletion argument (certified margin ≥ 0.046). Both confirm the same qualitative conclusion through independent analytical routes and are available from the authors upon request. The exact asymptotic identity $\mu_E \phi - S_G(\tau) \approx 0.216$ stated above is tighter than either alternative bound.

with Max Efficiency recovering later and purchasing for slightly longer. What matters is not who is ahead at every instant, but whether the early selection gains are large enough to survive the later recovery.




Figure 12: Logic of the asymptotic dominance result. The top panel shows the average ecological value of the plot purchased at each normalised turn under Max Environment and Max Efficiency in the baseline Budget World benchmark. The curves cross, so the theorem is not a pointwise-dominance result: Max Efficiency recovers later and also purchases for slightly longer. The bottom panel integrates those per-turn differences into the cumulative Purchased Conservation gap, $(PC^{ME} - PC^{MX})/N$. The contested phase (during which the Farmer is still shopping) creates a large advantage for Max Environment that the post-Farmer phase cannot undo, which is the economic content of the asymptotic dominance result. Vertical markers: prefix share $m = 0.10$, Farmer exhaustion $\bar{\tau} \approx 0.305$ in the fluid limit. Baseline $\rho = 0$, $N = 2000$, 3000 replications.

The bottom panel confirms that they are: the cumulative gap builds rapidly through the contested phase, and the late-game recovery by ratio-greedy never erases it. The asymptotic dominance result is therefore an integral statement about dynamic selection quality: value-first play wins because its contested-phase gains in ecological quality are larger than the value that cost-effectiveness recovers later, not because it dominates in every period.

When does the knapsack reversal hold? The ecological selection premium. The mechanism described above — ratio-greedy retreating into a corridor of ecologically weaker

sites, value-first competing for the best remaining ecology — requires that the two strategies actually select measurably different portfolios. When ecological values are nearly homogeneous, they do not: the top- e and top- e/a rules select similar sites and the reversal disappears. To characterise this requirement, we define a scalar diagnostic for valuation heterogeneity, the *Valuation Selection Premium* (VSP): the percentage by which the average focal-agent valuation of the top-valuation portfolio exceeds that of the top-ratio portfolio. In the conservation application studied here, the focal valuation is ecological value, and VSP specialises to the *ecological selection premium* (ESP):

$$\text{ESP} = \frac{\bar{E}_{\text{top-}e} - \bar{E}_{\text{top-}e/a}}{\mu_e}.$$

When ESP is large, value-first has a meaningful quality advantage to exploit in the prefix phase; when it is small, the two rules select similar portfolios and the knapsack reversal disappears. Compressed- E simulations show that the exact reversal holds whenever ESP exceeds roughly 12.5% of mean e , while the current proof architecture requires approximately 15%. On the paper’s baseline grids, ESP is about 53%; on the Bolivian landscape, it exceeds 30%. Both are well above either threshold. ESP is straightforward to compute from any landscape dataset, offering practitioners a diagnostic for whether the conditions for a knapsack reversal are likely to hold.

Formal results. The economic narrative above is formalised in two asymptotic dominance theorems.⁵ Under the paper’s baseline data-generating process, value-first asymptotically dominates ratio-greedy on Purchased Conservation:

Theorem 4.1 (Asymptotic dominance under independence). *Consider Budget World with budget parity, a naïve Farmer who moves first, full leakage ($L = 1$),⁶ and i.i.d. draws $(E_i, A_i) \sim U[0.1, 10]^2$ with $E_i \perp A_i$. Let ϕ be the MX Farmer-exhaustion share, defined by $\int_{1-\phi}^1 Q_A(v) dv = \mu_A/2$, and let τ be the ME Farmer-exhaustion share, defined by $S_F(\tau) = \mu_A/2$ in the fluid limit.*

⁵A plain-language guide to the economic content of the proofs is available as a companion document: Weinhold (2026), “The Economics Behind the Proofs.”

⁶Both theorems formalize Budget World with full leakage ($L = 1$): Green expenditure does not reduce Farmer’s remaining budget, and both teams spend until exhaustion. Under general $L \in [0, 1)$ the Farmer’s residual budget is burned on contested purchases by $(1 - L)a_i$ per capture, which couples the two budgets dynamically and breaks the clean Farmer-exhaustion crossing time that the fluid-limit proofs rely on. A rigorous $L < 1$ extension would require specifying the integer burn process, rounding rules, and move-order timing carefully; we do not pursue that here. Budget World simulations across $L \in \{0, 0.25, 0.5, 0.75, 1\}$ (Section 4.1) indicate the knapsack reversal persists empirically.

Then the asymptotic gap admits the closed-form identity

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left[\text{PC}_N^{\text{ME}} - \text{PC}_N^{\text{MX}} \right] = \mu_E \phi - S_G(\tau) \approx 0.2161 > 0.$$

Proof roadmap. Appendix A8.3 gives the proof overview. The baseline proof, which exploits an exact E–A symmetry of the fluid-limit dynamics under independence, is in Appendix A8.4, with supporting inputs in Appendix A8.5. \square

The baseline is an exact asymptotic identity; the $\rho = 0.3$ extension is a certified lower bound, because the $\rho = 0.3$ fluid-limit ODE lacks the closed-form symmetry that makes the baseline identity available. Both results are proved using a unified fluid-limit architecture that tracks Green’s actual per-turn ecological value through a deterministic ODE. Exact-object simulations at moderate N confirm gaps of roughly 0.216 at baseline and 0.29 at $\rho = 0.3$.

Theorem 4.2 (Asymptotic dominance at $\rho = 0.3$). *Under the same conditions but with (E_i, A_i) drawn from a Gaussian copula with latent correlation $\rho = 0.3$ and uniform marginals on $[0.1, 10]$,*

$$\liminf_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left[\text{PC}_N^{\text{ME}} - \text{PC}_N^{\text{MX}} \right] \geq 0.19 > 0.$$

Proof roadmap. Appendix A8.3 summarizes the correlated extension and records the constant ledger. The full $\rho = 0.3$ proof, including the validated Picard step-tube certificate for the fluid-limit ODE, is in Appendix A8.6 (with the certificate subsection Appendix A8.6.7). The certified margin is ≥ 0.198653 ; the baseline theorem (Theorem 4.1) closes independently through a closed-form symmetry identity and does not depend on the certificate subsection. \square

Honest boundaries. Universal dominance does not hold for the finite discrete game. Constructive counterexamples in the broader finite equal-budget procurement game establish that on very small grids, where a single asset can exhaust the budget, ratio-greedy can outperform value-first:

Proposition 4.3 (No universal finite-instance dominance in the equal-budget procurement game). *In the finite discrete equal-budget procurement game, even with naïve Farmers, equal budgets, full observability, and full leakage, neither MAX ENVIRONMENT nor MAX EFFICIENCY weakly dominates the other on Purchased Conservation over all instances.*

Proof. Appendix [A8.2](#) gives two constructive examples. □

These counterexamples rely on tight-budget mechanisms that become less representative under the benchmark sampling scheme as N grows: the budget concentration ratio $\kappa(N)$, which measures the share of the Green budget consumed by the single most expensive plot, shrinks like $O_p(1/N)$ under the baseline draws (Appendix [A8.8](#)). This does not by itself prove finite- N dominance of Max Environment, but it helps explain why the tiny-grid reversals disappear in the larger benchmark simulations. By 3×3 the ranking already aligns with the Monte Carlo pattern, and the Bolivian case study ($N = 390$, $\kappa \approx 0.01$) sits deep in the region where self-hedging dominates.

5 Empirical Case Study: Conservation Strategies in Bolivia

The simulations and formal theory above are built on stylised grids with i.i.d. draws and parametric copula structure. Robustness checks in Appendices [A2–A6](#) confirm that the qualitative rankings survive individual departures from the baseline, but an important question is whether they survive on a real landscape where multiple departures converge: environmental values may be heavy-tailed, agricultural values follow soil and transport gradients, the two dimensions are positively correlated through natural processes rather than by construction, and spatial clustering can concentrate high-value sites in particular regions. Bolivia provides a suitable test case: high-resolution biophysical data yield an estimated ecological–agricultural correlation of $\rho \approx 0.29$, close to the $\rho = 0.3$ benchmark for which the stylised knapsack reversal is formally proved, and the ecological selection premium on the Bolivian board exceeds 30%, well above the diagnostic range identified in the compressed- E exercises.

With extraordinary biodiversity and significant pressure for economic development, Bolivia’s history of conservation efforts also illustrates the adversarial dynamics that the framework formalises. Starting with the first national park in 1939, the land area under Protected Area (PA) status in Bolivia has gradually increased to at least 38.6 million hectares today, corresponding to more than one third of the national territory ([SDSN Bolivia, 2025](#)). 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, 2025). 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 14 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 14 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 13: Bolivian GDP per capita, 1950–2019
(constant 2017 USD adjusted for cost of living)

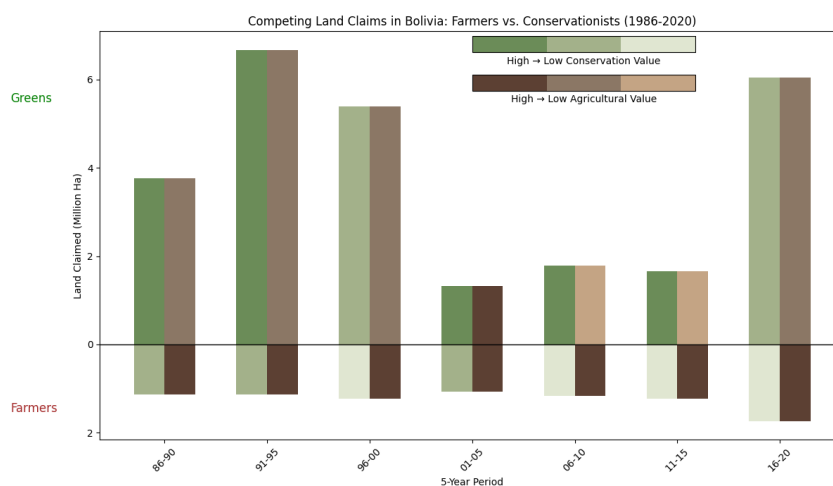


Figure 14: Agricultural & Conservation Potentials and Area of New Land allocated to Agriculture and Conservation, 5-year increments

Comparing the potential agricultural values (shades of brown) in Figure 14 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 13 shows the evolution of GDP per capita over the same period. Comparing this with Figure 14 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 15(a), overlaid 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. (2025). 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 $\rho = 0.29$, close to the $\rho = 0.3$ positive-correlation case analysed in Appendix A8.6. The ecological selection premium on this board exceeds 30%, well above the 12–15% threshold identified in Section 4.1.1, placing Bolivia squarely in the region where the knapsack reversal operates. The distribution of agricultural and environmental values is displayed in Figures 15(b) and 15(c).

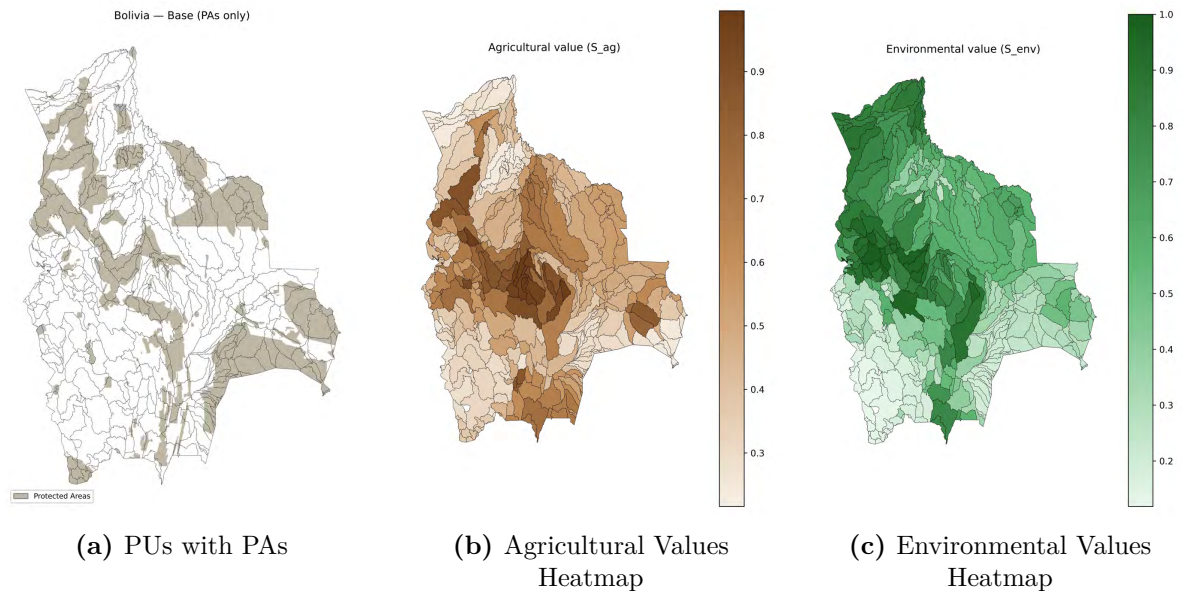


Figure 15: Bolivian PU board with 2024 Protected Areas and Heatmaps

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 16 - 19 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.

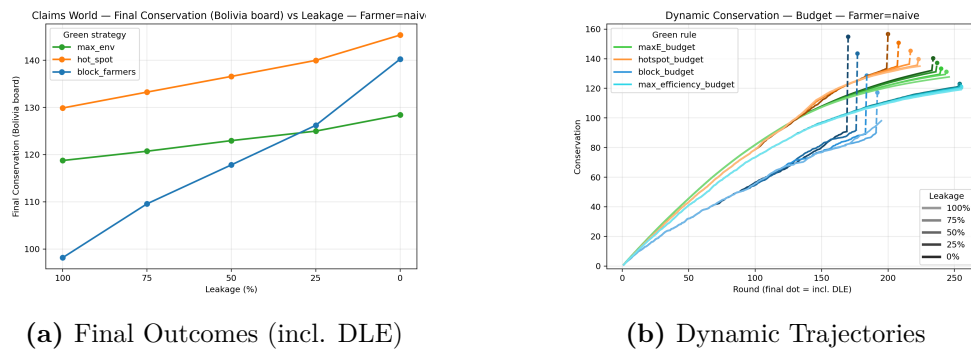


Figure 16: Claims World on a Bolivian Board with Naïve Farmers: Conservation Scores

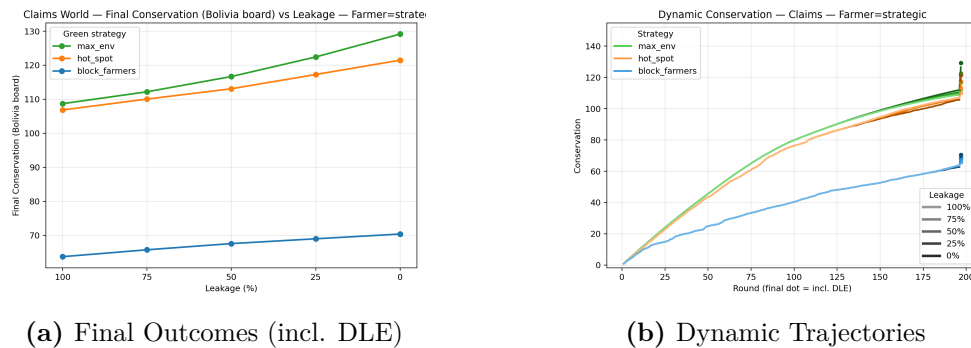


Figure 17: Claims World on a Bolivian Board with Strategic Farmers: Conservation Scores

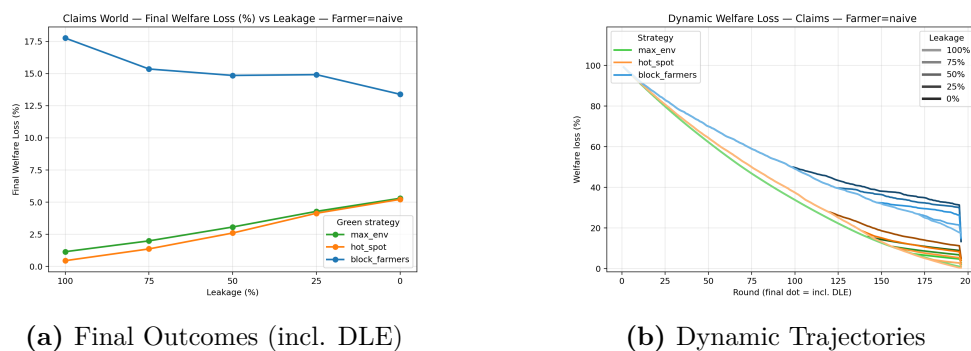


Figure 18: Claims World on a Bolivian Board with Naïve Farmers: Welfare Loss (%)

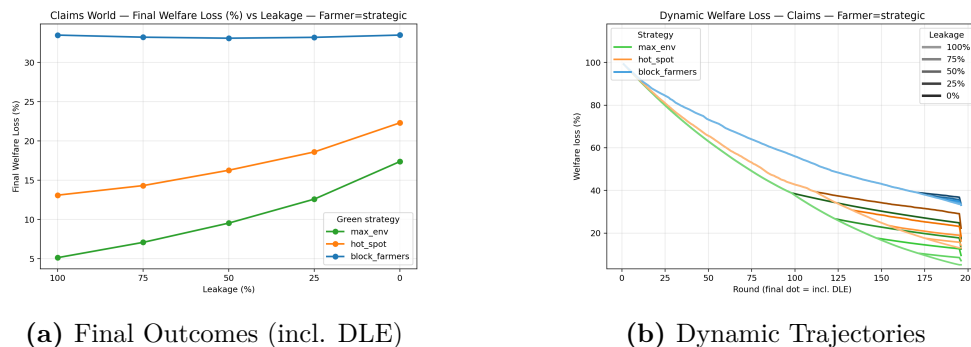


Figure 19: Claims World on a Bolivian Board with Strategic Farmers: Welfare Loss (%)

Claims World conservation outcomes on the Bolivia board show the same qualitative ranking of Green strategies we observe from the more stylized Monte Carlo (Figures 1 - 2 and Figures 5 - 6). Against Naïve Farmers (Figure 16), 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 16(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 17) 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 17(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 18 and 19 show welfare losses on the Bolivia board that mirror the patterns from the Monte Carlo (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 20 - 21) likewise confirm the robustness of the main simulation findings (the remaining results on additivity are similarly consistent with the Monte Carlos but not presented here as the patterns are identical to those for Conservation).

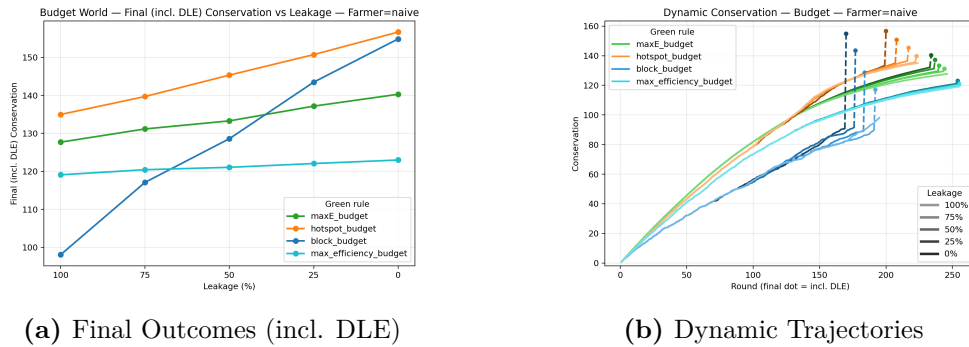


Figure 20: Budget World on a Bolivian Board with Naïve Farmers: Conservation Scores

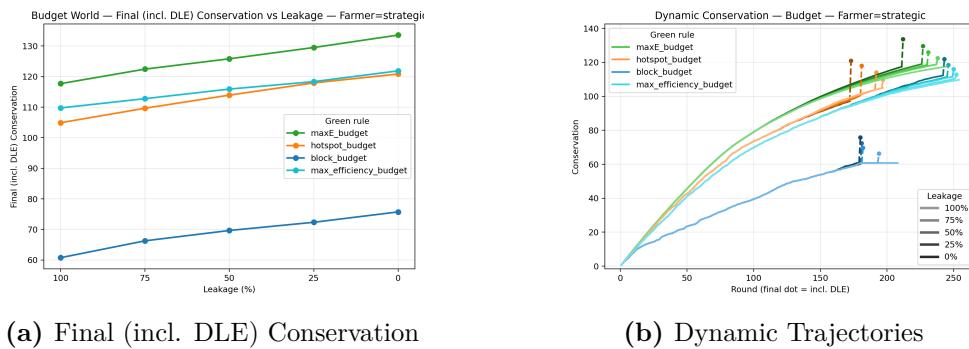


Figure 21: Budget World on a Bolivian Board with Strategic Farmers: Conservation Scores

In Budget World, against Naïve Farmers, Hotspot_budget achieves the highest final conservation at most leakage levels, with MaxE_budget close behind; both substantially outper-

form 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 gap that preserves the dominance ordering at every leakage level, reproducing on the Bolivia board the same selection effect analysed in Appendix A8.6.

Figures 22 and 23 show the spatial distribution of the chosen PUs by MaxE_budget and Hot Spot_budget against both Naïve and Strategic Farmers.

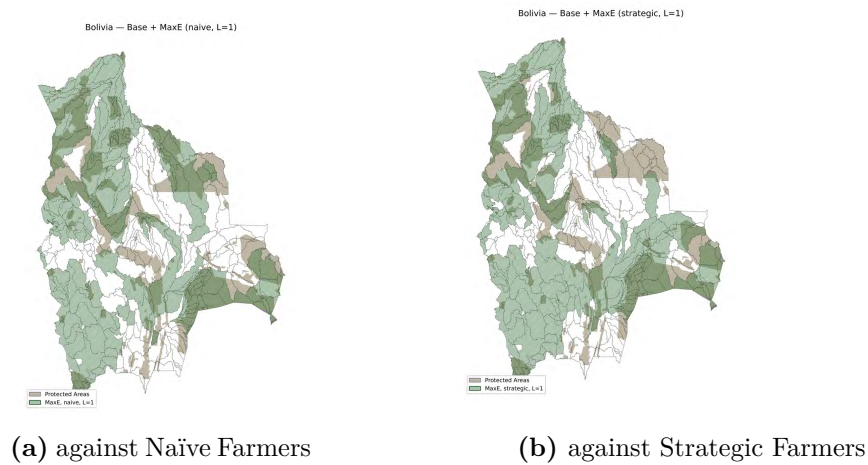


Figure 22: Budget World Max Environment Strategy Conserved PUs (leakage=100%)

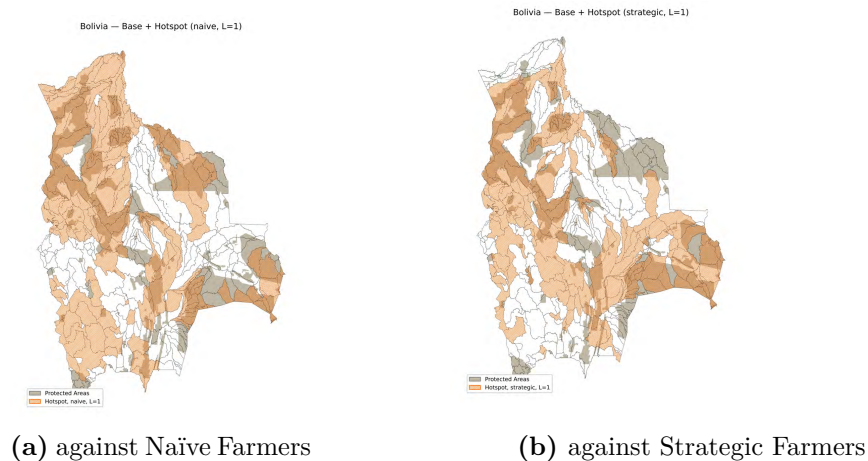


Figure 23: Budget World Hot Spot Strategy Conserved PUs (leakage=100%)

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 (?). 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. In the paper’s adversarial simulations, that heuristic performs poorly when costs also signal rival demand. To examine whether the same qualitative pattern appears 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 introduce the Marxan style strategy to the Budget World contest on the fixed Bolivia board. The results, presented in Figures 24 and mapped out in Figure 25, yield two striking and related insights regarding the performance of Marxan conservation strategies in a real-world contested environment.

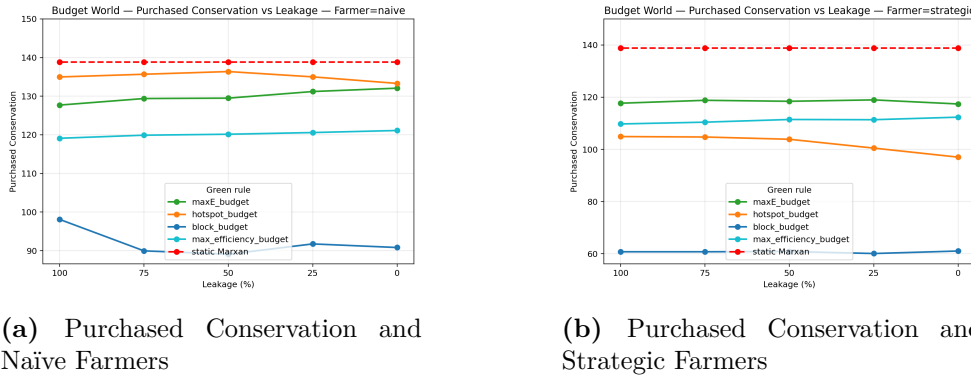


Figure 24: Budget World on a Bolivian Board: Final Purchased Conservation

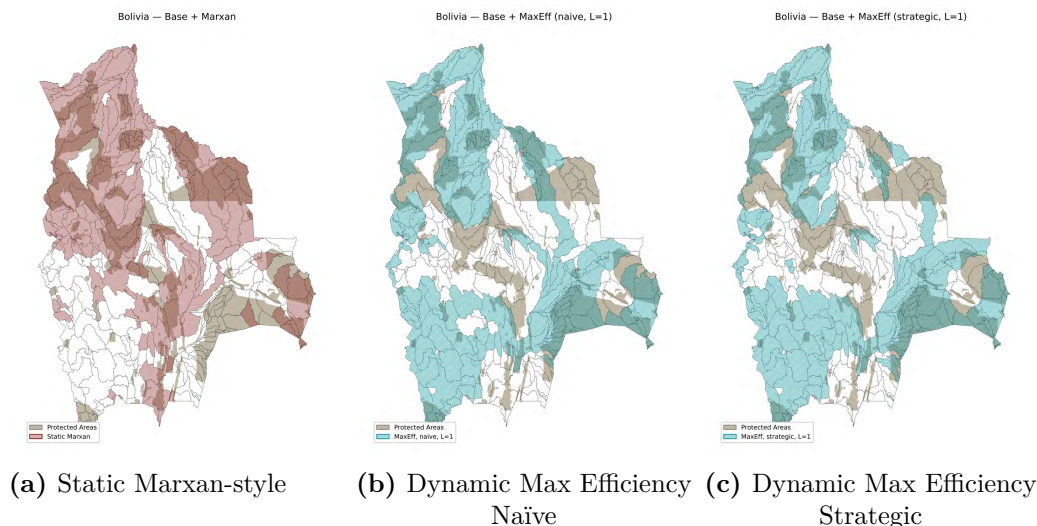


Figure 25: Static Marxan-style Benchmark and Dynamic Max Efficiency Conserved PUs (Budget World, leakage=100%)

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). The mechanism is the same self-hedging logic that drives the knapsack reversal: the static ratio-greedy benchmark favours cheap, high-ratio sites in the low- a corridor that the rival has little reason to target, while the truly contested sites slip away. The static plan therefore takes credit for conserving sites that were never at risk.

Second, the simulations show that the qualitative ranking from the stylized models survives on this real-world planning board. The simple **Max Environment** strategy (green line) consistently outperforms the **Max Efficiency** strategy (cyan line) in Purchased Conservation (PC) on the reported exercises. Bolivia’s estimated correlation ($\rho \approx 0.29$) falls close to the positive-correlation benchmark $\rho = 0.3$ for which the stylised knapsack reversal is formally proved in Theorem 4.2, and the board’s ecological selection premium exceeds 30% — well above the diagnostic range identified in the compressed- E exercises — so the dominance of value-first over ratio-greedy on this landscape is consistent with the formal result, not merely a simulation regularity.

6 Discussion

This paper identifies a structural class of allocation problems that we call *adversarial procurement in two-value space*: contests in which a budget-constrained agent acquires assets sequentially from a shared and shrinking pool, competes against an active rival, values assets differently from that rival, and pays prices determined by the rival’s valuation rather than its own. The conservation problem studied here is one important instance of that class, but the logic is broader. We study this structure computationally, formally, and on a real landscape, converging on three nested lessons: rivalry changes the meaning of threat, two-value space changes the meaning of cost, and valuation heterogeneity determines how much those distortions matter.

The Claims World analysis isolates the first lesson. Dynamic rivalry is the engine: threat is not a fixed property of a site but an endogenous object that changes as conservation and development respond to one another. The PSE/DLE decomposition makes that logic transparent. Through the *Pure Strategy Effect*, strategies that focus on the most threatened sites neglect intrinsically valuable sites that developers remove later, generating a loss that accumulates regardless of leakage. Through the *Displacement–Leakage Effect*, blocking a developer’s preferred site may or may not reduce total development, depending on how easily development re-targets. When leakage is high, DLE evaporates because the developer simply moves to the next best alternative. When leakage is low, DLE can be substantial, but even then threat-targeting is fragile because it depends on a particular developer response. The robust feature of Max Environment in Claims World is therefore not that it somehow blocks development more cleverly, but that it secures intrinsically valuable sites regardless of whether developers are naïve or strategic.

For conservation, a direct implication of this mechanism concerns the current emphasis on threat-based additionality in voluntary carbon markets and payments for ecosystem services. Additionality metrics are appealing because they aim to prioritise places where protection changes the counterfactual most. But in the adversarial settings studied here, measured additionality and realised conservation outcomes are not the same objective. Threat-based targeting channels resources toward contested frontiers where DLE may look large *ex ante* but can evaporate once developers re-target, while the PSE from neglecting intrinsically valuable

sites remains. The welfare paradox under reduced leakage sharpens that point: policies that reduce re-targeting can raise welfare losses if they end up steering high-agricultural-value land into conservation without commensurate ecological gain. In practice, this means policy should attend to the ecological quality of the sites developers are displaced toward, not only to the quantity of avoided conversion — though credible estimates of displacement patterns remain scarce and geographically local (Searchinger et al., 2018; Silva and Nunes, 2025).

Budget World reveals the second lesson and the paper’s most striking result. Once the focal agent’s acquisition cost is determined by the rival’s valuation rather than its own, cheapness is not safety-neutral information. The mechanism we call *self-hedging* operates: the same agricultural value that makes a plot expensive also makes it likely to be removed by the rival. In that setting, the static logic of the knapsack problem is reversed. A ratio-greedy rule saves budget by stepping away from the assets it can least afford to lose. The formal theory shows that this is not a simulation artefact. Theorem 4.1 proves asymptotic dominance of Max Environment over exact ratio-greedy Max Efficiency in the baseline benchmark, and Theorem 4.2 proves the same dominance at $\rho = 0.3$, the positive correlation estimated on the Bolivian landscape. The direct comparison in Section 4.1.1 makes the economic logic clear: value-first play builds an early Purchased Conservation lead, retains and often enlarges that lead while the rival is still active, and then gives some of it back late in the game, but not enough to reverse the ranking. The theorem is thus not a pointwise-dominance claim; it is an integral statement about dynamic selection quality. The late-game recovery by the ratio rule is real, but it is too small to erase the early advantage. More generally, the economics of the knapsack reversal lives primarily in early selection quality, not in late-game budget management: late dynamics matter because they test whether the early advantage survives, but they do not create that advantage.

The Budget World diagnostics clarify the role of Farmer behaviour. The formal theorems are proved for a max- A Farmer because that case makes the self-hedging channel analytically visible: Farmer interference removes expensive items, and the resulting budget relief can be controlled in the fluid limit. The simulations suggest the knapsack reversal is more general. Max Environment continues to dominate Max Efficiency against Strategic Farmers, whose BAU set is explicitly shifted toward plots that Green could credibly acquire. The broader mechanism at work is the same self-hedging logic: in adversarial procurement, A is not only Green’s acquisition price but also the rival’s value and, under budget parity, the rival’s cost of denial. A high- e , high- a plot

is therefore not simply expensive — it is contested. Value-first forces the rival either to let Green obtain ecological value or to spend scarce capacity denying it; ratio-greedy may decline the same plot on cost-effectiveness grounds, allowing the rival to capture it without being made to pay.

The leakage sweeps suggest this broader pattern operates through multiple reinforcing channels. At high leakage, Strategic Farmers create a *contestation premium*: because their BAU set contains more ecologically valuable plots, delaying high- e acquisitions is especially costly. At low leakage, a *blocking channel* becomes more important, because Green's purchases of Farmer-BAU plots reduce future Farmer capacity. The absolute magnitude of the leakage-driven rise in the Naïve-Farmer gap is nearly identical at $\rho = 0$ (+10.2 units) and $\rho = 0.3$ (+10.0 units), suggesting that the blocking channel is approximately correlation-invariant even though the selection-quality floor on which all channels operate is amplified by correlation. The reversal therefore does not rely on any single behavioural assumption; the present theory proves only the clean max- A case, and extending the formal result to strategic budgeted rivals is a natural next step.

The comparative statics yield a surprising insight: positive correlation between ecological and agricultural value *amplifies* the knapsack reversal rather than attenuating it. At $\rho = 0.3$, the reversal is formally proved, and the computational sweep in Figure 11 shows that the normalised Purchased Conservation gap remains positive throughout the interior positive-correlation range we examine and is strongest at intermediate positive correlation. The economic reason is that partial alignment makes the ecologically best plots more contested, while the ratio-greedy rule continues to favour sites with anomalously low agricultural value relative to ecological value. The knapsack reversal weakens only as ρ approaches one, when the two dimensions become nearly collinear and the distinction between value-first and ratio-first selection largely disappears. The implication is that the danger of cost-effectiveness mandates is greatest not when ecological and agricultural values are unrelated, and not when they are identical, but when they are substantially and imperfectly aligned.

The compressed- E exercises identify the other side of the scope condition. The knapsack reversal weakens when ecological values become too homogeneous, not because the budget or stopping-time mechanics change very much, but because the ecological quality premium from

value-first selection disappears. This motivates a scalar diagnostic for valuation heterogeneity, which we call the *Valuation Selection Premium* (VSP) in the general setting and, in the conservation application studied here, the *ecological selection premium* (ESP): the percentage by which the average ecological value of the top- e portfolio exceeds that of the top- e/a portfolio. When ESP is large, Max Environment has a meaningful quality advantage to exploit; when ESP is small, the two rules select increasingly similar portfolios and the reversal vanishes. In the compressed- E diagnostics, the exact knapsack reversal persists down to an ESP of roughly 12.5% of mean e , while the current proof architecture certifies it down to about 15%. On the paper’s baseline grids, the premium is about 53%, and on the Bolivian landscape it exceeds 30%. This gives practitioners a useful diagnostic: before adopting a cost-effectiveness rule, ask whether ecological heterogeneity is large enough for value-first selection to buy genuinely better land.

The theorems are proved under budget parity, but the phase decomposition suggests that parity is the conservative benchmark for value-first play rather than a case specially favourable to it. A larger Farmer budget extends the contested phase of the game — precisely the phase in which Max Environment gains ground — and compresses the late uncontested phase where Max Efficiency partially recovers. In real-world conservation, the agricultural sector’s effective “budget” is often larger than the conservation budget, so the contested region that drives the self-hedging advantage is likely longer, not shorter, than in the equal-budget benchmark. The unequal-power simulations in Appendix A4, which allocate 70% of claims to Farmers, are consistent with this interpretation: the qualitative ranking of strategies is preserved.

Bolivia provides a direct empirical test on a realistic and highly heterogeneous planning board. With spatially correlated values, heavy-tailed environmental heterogeneity, and an estimated positive ecological–agricultural correlation of $\rho \approx 0.29$, the Bolivian landscape falls close to the positive-correlation benchmark $\rho = 0.3$ for which the stylised knapsack reversal is formally proved — and the qualitative rankings match the theory exactly. More practically, the Bolivian exercises show that static Marxan-style cost-effectiveness plans overestimate realised conservation value by roughly 15–20% once developer responses are allowed for. That gap is large enough to matter for conservation finance, carbon-credit methodologies, and other policy settings in which cost-effectiveness is often treated as the binding evaluation criterion.

Beyond conservation, the economic mechanisms our analysis reveals extend to any setting where a mission-driven buyer deploys a budget against a rival whose valuation helps set prices — the health systems, government procurement, and public land acquisition settings described in the introduction all share this two-value structure. The specific magnitudes will differ across domains, but the core mechanism is general: when acquisition cost tracks rival demand, the benefit-cost ratio's apparent savings can reflect uncontested assets the rival was never going to claim. In any such setting, decision-makers who understand how rivalry reshapes the meaning of cost will outperform those who optimise benefit-cost ratios derived from a static menu.

Taken together, the three nested lessons reveal adversarial rivalry in two-value space as a strategic environment with its own logic. Rivalry makes threat endogenous; two-value structure makes cost informative about contest; valuation heterogeneity determines the magnitude of both effects. Each lesson carries distinct empirical conditions and distinct policy implications, but all three interact: it is the *combination* of dynamic rivalry, multi-dimensional valuation, and heterogeneous focal-agent values that generates the paper's results. Where these conditions are met, the classical toolkit built for a static single-agent menu requires revision.

For conservationists, the practical lesson is not to follow any fixed rule rigidly, but to understand how rivalry and two-value structure reshape the forces at work. In rival-aware procurement, cheapness can be a warning sign rather than a margin of safety, and high measured additionality can signal a contested frontier where displacement-based gains are fragile. Taken together, our results suggest a robust decision rule: start from ecological value, then use cost and threat to refine that choice rather than to replace it.

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.
- Amar, J. and Renegar, J. (2023). The second-price knapsack problem: Near-optimal real-time bidding in internet advertisement. *arXiv preprint*.
- 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., Argandoña, F., Calderón, D., Choque, S., Muñoz, A., Olmos, C., and Miranda, S. (2025). Valoración económica de los servicios ecosistémicos en bolivia: estimación del aporte de las áreas protegidas y los territorios indígenas. *Latin American Journal of Economic Development*, 44:109–153.
- 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. (1998). 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.
- 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.
- Baum-Snow, N. and Marion, J. (2009). The effects of low income housing tax credit developments on neighborhoods. *Journal of Urban Economics*, 65(1):1–15.
- Brealey, R. A., Myers, S. C., and Allen, F. (2020). *Principles of Corporate Finance*. McGraw-Hill, New York, 13th edition.
- Colyvan, M., Justus, J., and Regan, H. M. (2011). The conservation game. *Biological Conservation*, 144(4):1246–1253.
- Costello, C. and Polasky, S. (2004). Dynamic reserve site selection. *Resource and Energy Economics*, 26(2):157–174.
- Dantzig, G. B. (1957). Discrete-variable extremum problems. *Operations Research*, 5(2):266–288.
- 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.
- Ethier, S. N. and Kurtz, T. G. (1986). *Markov Processes: Characterization and Convergence*. Wiley, New York.
- 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.
- 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.
- 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.
- Gold, M. R., Siegel, J. E., Russell, L. B., and Weinstein, M. C. (1996). *Cost-Effectiveness in Health and Medicine*. Oxford University Press, New York.
- Gummadi, R., Key, P., and Proutiere, A. (2012). Repeated auctions under budget constraints: Optimal bidding strategies and equilibria. In *Proceedings of the 8th Workshop on Ad Auctions (AdAuctions 2012)*.

- 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.
- Jaakkola, N. and Wagener, F. (2025). Markov-perfect equilibria in differential games, with an application to climate policy. *CESifo Working Papers*, (10585). Original version July 2023, revised October 2025.
- Johnson, M. P. and Hurter, A. P. (2000). Decision support for a housing mobility program using a multiobjective optimization model. *Management Science*, 46(12):1569–1584.
- 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.
- Kurtz, T. G. (1970). Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7(1):49–58.
- 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.
- Lueck, D. and Michael, J. A. (2003). Preemptive habitat destruction under the endangered species act. *Journal of Law and Economics*, 46(1):27–60.
- MapBiomass 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.

- Miceli, T. J. and Segerson, K. (2012). Land assembly and the holdout problem under sequential bargaining. *American Law and Economics Review*, 14(2):372–390.
- 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>.
- Neumann, P. J., Sanders, G. D., Russell, L. B., Siegel, J. E., and Ganiats, T. G. (2016). *Cost-Effectiveness in Health and Medicine*. Oxford University Press, New York, 2nd edition.
- 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.
- Pferschy, U., Nicosia, G., Pacifici, A., and Schauer, J. (2021). On the stackelberg knapsack game. *European Journal of Operational Research*, 291(1):18–31.
- Pollard, D. (1984). *Convergence of Stochastic Processes*. Springer, New York.
- SDSN Bolivia (2025). Áreas protegidas de bolivia, 2025. SDSN Bolivia - UPB, Datos Espaciales. Available at: <https://sdsnbolivia.org/datos-espaciales/>.
- Searchinger, T., Wiersenius, S., Beringer, T., and Dumas, P. (2018). Assessing the efficiency of land use changes for climate mitigation. *Nature*, 564(7735):249–253.
- Shavell, S. (2010). Eminent domain versus government purchase of land given imperfect information about owners' valuations. *Journal of Law and Economics*, 53(1):1–27.
- Silva, D. S. and Nunes, S. (2025). Leakage effects from reforestation: Estimating the impact of agricultural displacement for carbon markets. *Land*, 14(5):963.
- Thakoor, O., Kannan, R., and Prasanna, V. (2025). Adversarial knapsack for sequential competitive resource allocation. arXiv preprint.

- 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.
- Weinhold, D. (2026). The economics of adversarial procurement in two-value space: intuition behind the knapsack reversal. Available at https://github.com/dmweinhold/Latest-Version-Papers/raw/main/Bronze_theorem_economic_intuition.pdf.
- Weinstein, M. C. and Stason, W. B. (1977). Foundations of cost-effectiveness analysis for health and medical practices. *New England Journal of Medicine*, 296(13):716–721.
- Wormald, N. C. (1999). The differential equation method for random graph processes and greedy algorithms. In Karoński, M. and Rucinski, H.-J., editors, *Lectures on Approximation and Randomized Algorithms*, pages 73–155. PWN, Warsaw.
- Wu, J., Zilberman, D., and Babcock, B. A. (2001). Environmental and distributional impacts of conservation targeting strategies. *Journal of Environmental Economics and Management*, 41(3):333–350.
- Wunder, S. (2015). Revisiting the concept of payments for environmental services. *Ecological Economics*.

Appendix

This Appendix serves three purposes: it stress-tests the main simulation results against a wide range of alternative assumptions, it establishes the theoretical boundaries of the knapsack reversal, and it documents the data and institutional context for the Bolivian case study.

The robustness exercises vary modelling choices that might drive the results. Appendix A2 introduces positive correlation between agricultural and environmental values ($\rho = 0.3$), as suggested by the Bolivian data; Appendix A3 simulates weak legal enforcement via negative correlation ($\rho = -0.3$), capturing settings where protection is hardest to sustain on high-development-value land; and Appendix A4 tilts political power toward developers by allocating 70% of claims to Farmers. Appendix A5 replaces the baseline uniform distribution with heavy-tailed environmental draws, so that conservation outcomes hinge on a few “superstar” sites. Appendix A6 introduces spatial spillovers via randomly paired plots (S1), and Appendix A6.1 generalises this to a graph-based model (S2) that formalises why clustering favours value-first play. The qualitative strategy rankings survive all of these perturbations. Appendix A7 formally sets up the Budget World game.

The theory addendum (Appendix A8) establishes the formal structure of the knapsack reversal. Appendix A8.1 gives an exact Green-turn recursion for stripped Claims World under full leakage and explains why scalar threshold rules are insufficient in the full discrete game. Appendix A8.2 gives the finite equal-budget counterexamples and clarifies why no universal finite-instance dominance theorem is available. Appendix A8.3 provides the proof roadmap and the constant ledger for the asymptotic dominance theorems. Appendix A8.4 contains the full baseline proof, Appendix A8.5 collects the supporting inputs S1–S4, and Appendix A8.6 contains the direct $\rho = 0.3$ extension and its validated closure. Appendix A8.7 characterises the ecological conditions under which the knapsack reversal holds, Appendix A8.8 discusses finite- N budget concentration and small-grid pathologies, and Appendix A8.9 benchmarks all heuristics against backward-induction equilibria on 2×2 , 3×3 , and 4×4 grids. A brief companion guide to the intuition behind the theoretical proofs presented in Appendix A8 is provided in Weinhold (2026). Appendix A9 provides supplemental material on the history of protected areas in Bolivia and the construction of our high-resolution dataset of agricultural

and environmental values. Finally, to support replication and teaching, Appendix [A10](#) provides links to two publicly available tools: a standalone Python simulator (archived on Zenodo) for reproducing all reported results and exploring alternative parameterisations, and an interactive browser-based game where users can play as either team against computer-controlled strategies.

A1 Claims World Baseline Conservation Achieved

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

Farmer Strategy	Green Strategy	Leakage Level (%)	<i>Pure Strategy</i> Effect	<i>Displacement-Leakage</i> 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

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.

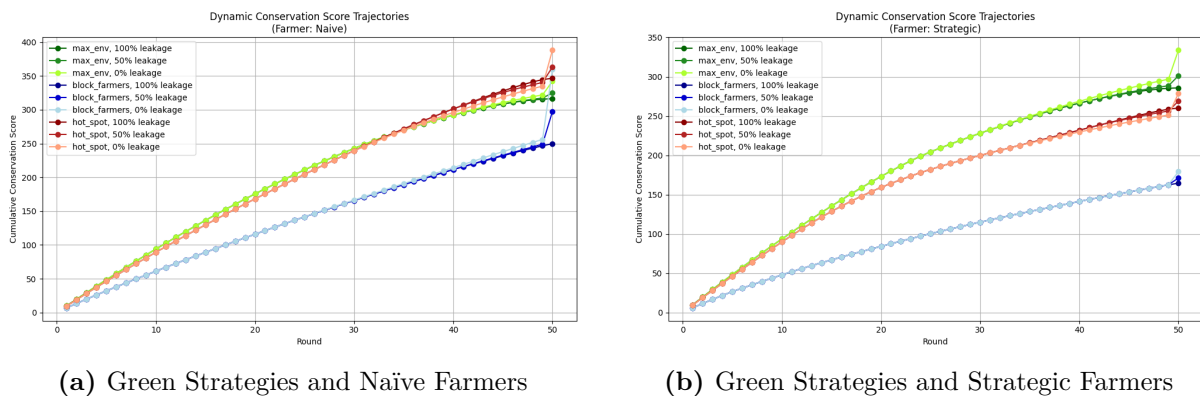


Figure A2.1: Conservation Scores for $\rho = 0.3$

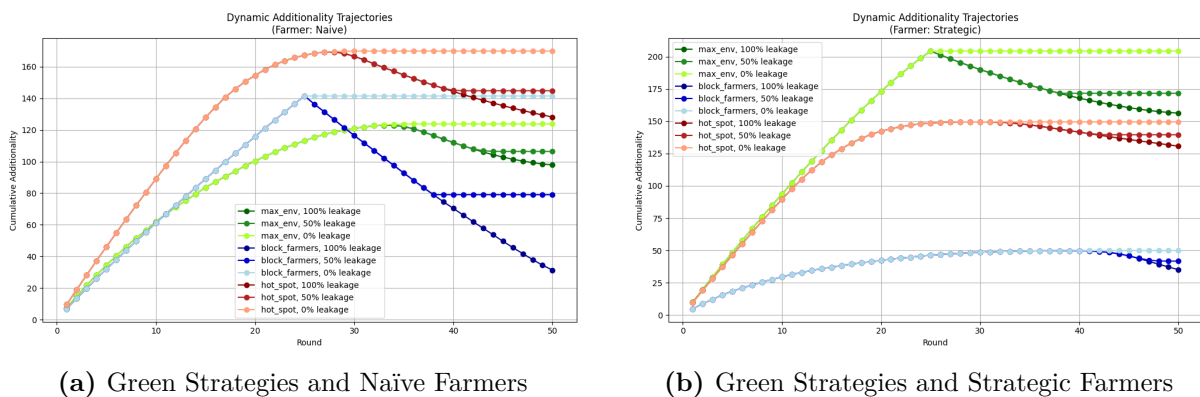


Figure A2.2: Additionality for for $\rho = 0.3$

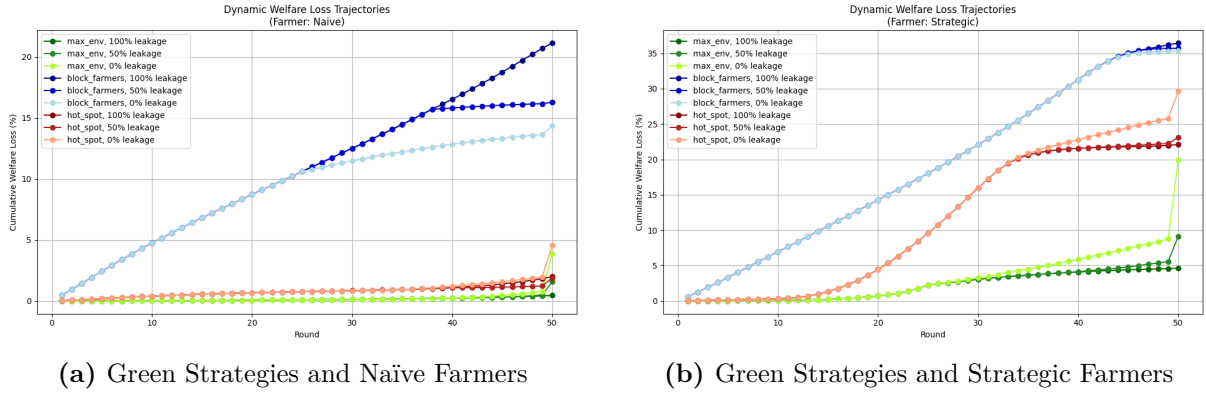


Figure A2.3: Social Welfare Loss(%) for $\rho = 0.3$


The positive-correlation simulation results reported above are now supported by formal theory. Theorem 4.2 in Appendix A8.3 proves that the knapsack reversal holds asymptotically at $\rho = 0.3$, with a certified margin of 0.198 per plot. The theoretical analysis reveals that positive correlation amplifies the reversal: the prefix wedge increases from 0.219 (baseline) to 0.309 ($\rho = 0.3$) because the high- E/A portfolio that Max Efficiency selects becomes more ecologically mediocre under positive correlation. Simulations indicate that the advantage is hump-shaped in ρ , peaking near $\rho \approx 0.7$ (Figure 11).

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 $\rho = -0.3$




(a) Green Strategies and Naïve Farmers




(b) Green Strategies and Strategic Farmers

Figure A3.2: Additionality for for $\rho = -0.3$



(a) Green Strategies and Naïve Farmers



(b) Green Strategies and Strategic Farmers

Figure A3.3: Social Welfare Loss(%) for for $\rho = -0.3$

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.

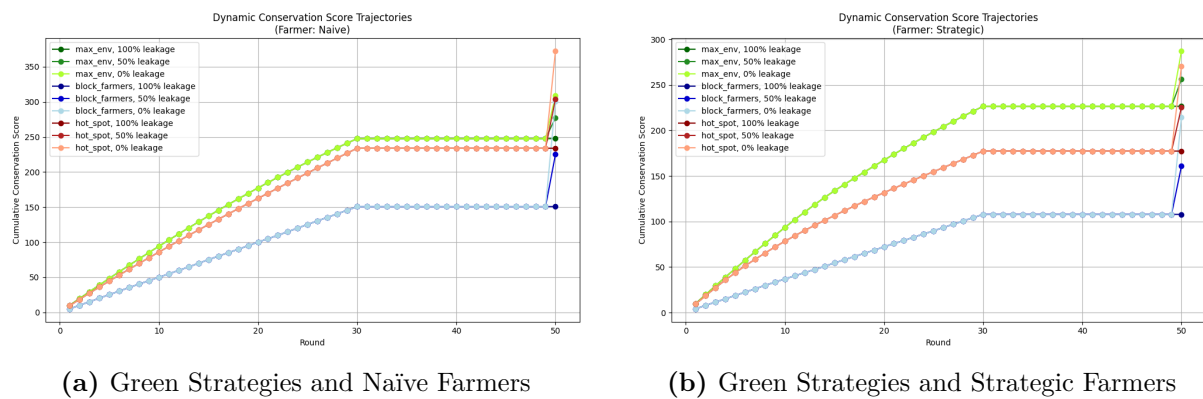


Figure A4.1: Conservation Scores for Political Allocation to Farmers = 70%

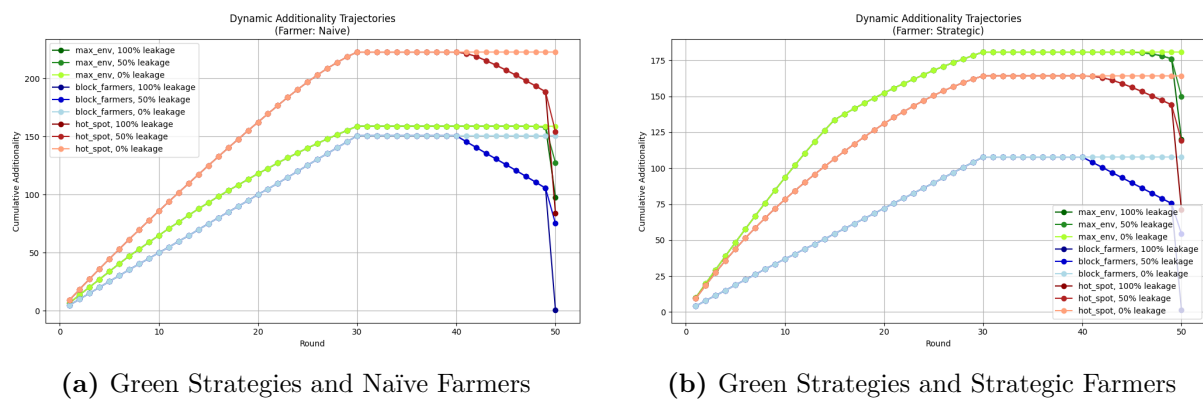


Figure A4.2: Additionality for Political Allocation to Farmers = 70%

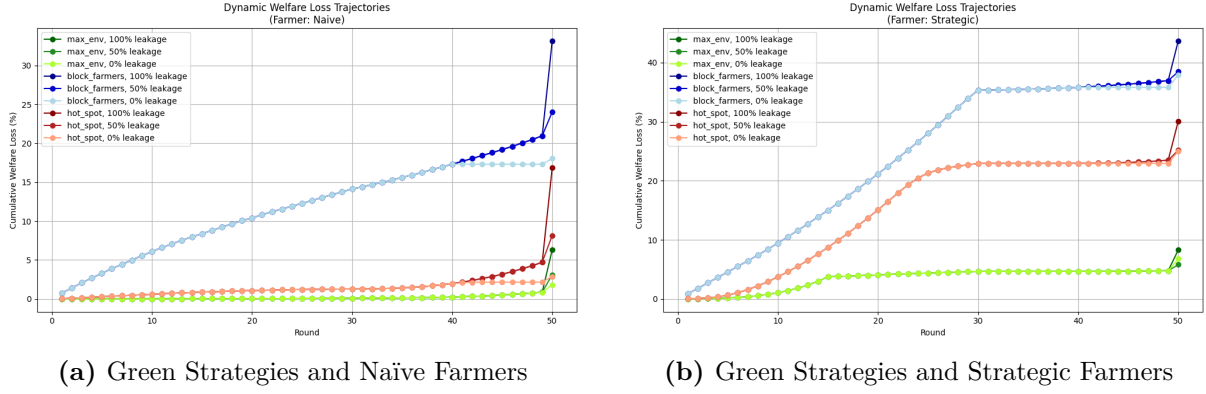


Figure A4.3: Social Welfare Loss(%) for Political Allocation to Farmers = 70%


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


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

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

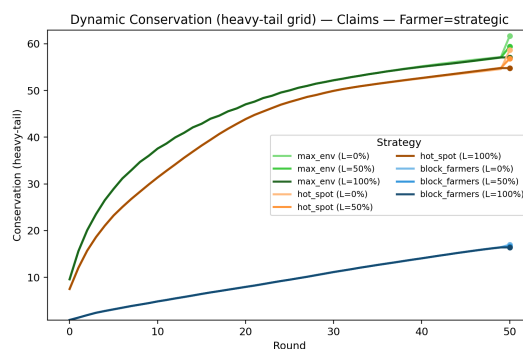


(b) Dynamic trajectories with heavy tailed e-values

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

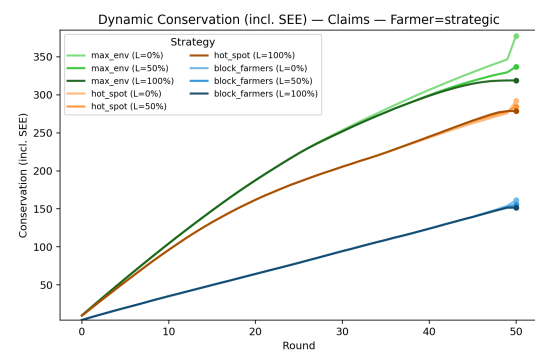


(b) Dynamic trajectories with spatial effects

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 Heuristic Implications for 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) discuss why, under mild assumptions, S2 may tilt play toward value-first conservation 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 $W_{ii} = 0$. We take W symmetric throughout (i.e. unnormalized, or row-normalized on a regular graph so that row-normalization preserves symmetry); the factor-of-2 cross-term in Proposition A6.2 below relies on this symmetry. Row-normalization on a non-regular graph yields $W_{ij} \neq W_{ji}$ in general, in which case the cross-term and marginal-gain formulas below must be replaced by their symmetrized analogues $\eta_G \sum_j (W_{ij} + W_{ji})$; the supermodularity conclusion carries through unchanged so long as the entries remain nonnegative.

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

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

with $\eta_G, \eta_F \geq 0$. The Green team’s directly claimed component of conservation (our PSE°) is

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

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

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

As in the body of the paper, the *Displacement–Leakage Effect* (DLE) credits residual unclaimed plots at their *base* ecological values e_i at the end of play. We define *Final Conservation* by the accounting decomposition

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

which mirrors the PSE/DLE convention in the main text while making spatial gains additive and visible.⁷

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. If S1 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 S1 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, $\sum_{j \in S_g} W_{ij}$ equals 1 if and only if $m(i) \in S_g$, and 0 otherwise; similarly for S_f . Substituting into (1) reproduces the S1 update rule. \square

Why S2 favors clustering

Define the set function on Green claims

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

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

Proposition A6.2 (Supermodularity; inclusion monotonicity). *Assume W is symmetric as*

⁷Equation (2) is an accounting convention rather than a physical identity: residual plots are credited at their base e , not at their spillover-adjusted e^{eff} , so that direct strategy effects stay visually distinct from mechanical residual gains. A spillover-adjusted variant would add the cross terms $\eta_G \sum_{i \in \text{Residual}} \sum_{j \in S_g} W_{ij} - \eta_F \sum_{i \in \text{Residual}} \sum_{j \in S_f} W_{ij}$; these terms do not affect the comparative statements below, which are made within the accounting convention (2).

specified in Section A6.1. If $W_{ij} \geq 0$ and $\eta_G \geq 0$, then $F(\cdot; S_f)$ is supermodular in S_g (increasing differences). The marginal gain from adding i to S_g is

$$\Delta_i(S_g; S_f) = e_i + 2\eta_G \sum_{j \in S_g} W_{ij} - \eta_F \sum_{j \in S_f} W_{ij},$$

and it is weakly increasing under set inclusion in S_g : if $S_g \subseteq S'_g$, then

$$\Delta_i(S_g; S_f) \leq \Delta_i(S'_g; S_f).$$

Moreover, $F(\cdot; S_f)$ is monotone in S_g provided

$$(4) \quad e_i \geq \eta_F \sum_{j \in S_f} W_{ij} \quad \text{for all } i \notin S_g \text{ and all feasible } S_f.$$

When $\eta_F = 0$ (conservation spillovers only), monotonicity holds unconditionally.

Proof. Symmetry of W gives $\sum_{j \in S_g} (W_{ij} + W_{ji}) = 2 \sum_{j \in S_g} W_{ij}$. Adding i to S_g therefore changes the Green–Green double sum by

$$2 \sum_{j \in S_g} W_{ij},$$

because both the new row and the new column contribute equally. Since $W_{ij} \geq 0$, this cross term is weakly increasing when S_g grows by set inclusion, which gives supermodularity.

Monotonicity requires $\Delta_i(S_g; S_f) \geq 0$, which holds under (4). □

The proposition formalizes a *connectivity dividend*: as a Green cluster grows, nearby plots become incrementally more attractive, independently of agricultural values a_i . This does *not* by itself prove that any particular fixed heuristic optimizes the spatial objective, but it does clarify the force introduced by S2.

Remark A6.3 (Heuristic implication for the strategy ranking under S2). Assume (A1) W is exogenous and independent of (e, a) ; (A2) $\text{corr}(e, a) = 0$; and (A3) condition (4) holds. Then S2 introduces a connectivity dividend that is aligned with ecological value and orthogonal to agricultural value. This makes it plausible that, among the fixed heuristics studied in the paper, MAX ENVIRONMENT should tend to benefit relative to a -weighted rules such as HOT SPOT and BLOCK FARMERS, especially as η_G and the average degree \bar{k} of W increase. We

treat this as a heuristic comparative-static interpretation rather than as a proved expected-payoff ranking theorem.

Remarks. (i) When leakage tends to one ($L \rightarrow 1$), the BAU displacement channel shuts down and the Green best response in the non-spatial game already shifts away from pure threat-chasing; S2’s spatial complementarity pushes further in the same direction. (ii) If the links in W are hidden from 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.⁸

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

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

so that the *per-neighbor* increment in e_i^{eff} matches the S1 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.

Take-away for the reader

S2 embeds a simple economic force: clustering by Greens creates value on neighboring plots (supermodularity), whereas weighting by agricultural threat does not by itself target that network dividend. Under (A1)–(A2), this makes it plausible that value-first heuristics retain or enlarge their advantage relative to a -weighted rules as spillovers strengthen. We present this as a heuristic robustness interpretation, not as a proved expected-payoff ranking theorem.

⁸See 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 as those that would be purchased by a Green Max Environment strategy operating under B_g^0 —i.e., the set of plots a greedy-by- e buyer could afford given the Green budget. Unlike the Claims World definition, which uses a simple rank cutoff (top- $\lceil \tau N \rceil$ by e), the Budget World risky set conditions on affordability, excluding high- e plots whose price a_i would prevent Green acquisition. Purchase risky plots in descending a_i ; once the risky set is exhausted, continue purchasing safe plots in descending a_i .

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, motivated by the fractional knapsack relaxation but generally not exact for the 0–1 problem.⁹

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.

⁹In 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: Formal Heuristic Comparisons and Budget-World Dominance

This appendix is self-contained. We first give exact limiting-case recursions for stripped Claims World objects (Section A8.1), then turn to Budget World, where the main results reside. Propositions A8.5 and A8.6 give constructive counterexamples showing that no universal finite-instance dominance theorem holds in the broader equal-budget procurement game. Theorems 4.1 and 4.2 then show that, under the paper’s baseline data-generating process, value-first asymptotically dominates ratio-greedy on Purchased Conservation at the baseline independence case and at the positive-correlation case $\rho = 0.3$.

Section A8.3 provides the proof roadmap and constant ledger. Section A8.4 contains the full baseline asymptotic dominance theorem proof, Section A8.5 collects the supporting input lemmas S1–S4, and Section A8.6 contains the direct $\rho = 0.3$ extension. The remaining subsections give ecological-condition diagnostics, budget-concentration results, and tiny-grid backward-induction benchmarks that clarify scope and finite- N behavior. A companion document (Weinhold, 2026, “The Economics Behind the Proofs”) provides plain-language economic intuition for each step of the proof architecture.

A brief companion guide to the intuition behind the theoretical proofs presented here (Weinhold, 2026) can be downloaded [here](#): [Companion Guide to Theory](#).

A8.1 Claims World: limiting-case intuition

The exact finite Claims World depends on the full remaining state. The results in this subsection are therefore deliberately narrow: they give an exact recursion under full leakage and then use remarks, not blanket theorems, to interpret the zero-leakage and strategic-Farmer cases.

Full leakage ($L = 1$).

Proposition A8.1 (Exact Green-turn recursion in the stripped alternating one-pick Claims World under full leakage). *Restrict Claims World to the following stripped game: one Farmer*

claim and one Green claim per round, Farmers move first, leakage is full ($L = 1$), Farmers play the naïve Profit Maximizer, and ties are resolved by a fixed measurable selector. Let $A \subseteq P$ denote the set of plots available at a Green turn, and let $j(S)$ denote the selected highest- a plot from any nonempty set S . Then Green's continuation value satisfies

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

with terminal condition $V_g(\emptyset) = 0$ and the convention that $j(\emptyset)$ is omitted.

Proof. At a Green turn with available set A , Green chooses some $i \in A$. Under full leakage the Farmer's continuation action is unaffected by burn, so the Farmer next removes the selected highest- a plot from $A \setminus \{i\}$. Under the stripped alternating one-pick timing, the game then returns to a Green turn on the smaller available set

$$A \setminus \{i, j(A \setminus \{i\})\}.$$

Maximizing over i gives the recursion. □

Remark A8.2 (Interpretation at full leakage). Proposition [A8.1](#) is intentionally narrow: it is an exact recursion for the stripped alternating one-pick, full-leakage Claims World, not for the more general multi-pick timing described in Section [2](#). In general the optimal choice depends on the full remaining set A ; a scalar comparison between the currently threatened plot and the best safe plot is not sufficient. On small grids the optimal action often resembles a preempt-or-value heuristic, but we do not claim an exact threshold rule or a subgame-perfect equilibrium characterization for the full discrete game.

Zero leakage ($L = 0$).

Remark A8.3 (Interpretation at zero leakage). When $L = 0$, capturing a Farmer BAU plot reduces the Farmer's future claim capacity. This creates a genuine blocking premium, but in the exact discrete game that premium depends on the full remaining state, not on a single scalar bonus θ_s . It is therefore more accurate to interpret Hot Spot-style behaviour as a heuristic response to a positive blocking premium than as an exact theorem.

Remark A8.4 (Strategic Farmers). If the Farmer's BAU set is loaded toward high- e plots,

delaying high- e purchases becomes more dangerous. This pushes the equilibrium logic toward value-first play, but again the exact continuation value is state dependent.

A8.2 Budget World: exact results and interpretation

We study Purchased Conservation (PC), which strips out the residual crediting channel through which leakage primarily affects final conservation. In the full dynamic game the purchased set can still vary with L via Farmer budget burn, but on the paper’s baseline grids that feedback is modest relative to the selection effect that drives the head-to-head comparison between MAX ENVIRONMENT and MAX EFFICIENCY.

A8.2.1 Baseline head-to-head simulation evidence

Simulation results (Section 4.1) show that MAX ENVIRONMENT strongly outperforms MAX EFFICIENCY on purchased conservation (PC) in the paper’s baseline Monte Carlo design across all reported leakage levels and against both naïve and Strategic Farmers. Figure A8.1 zooms in on that head-to-head comparison, and Table A8.1 reports the corresponding win-rates and PC gaps across 500 replications per leakage level for each Farmer type. We treat this as a robust empirical regularity under the benchmark design, not as a universal theorem for the exact finite game.

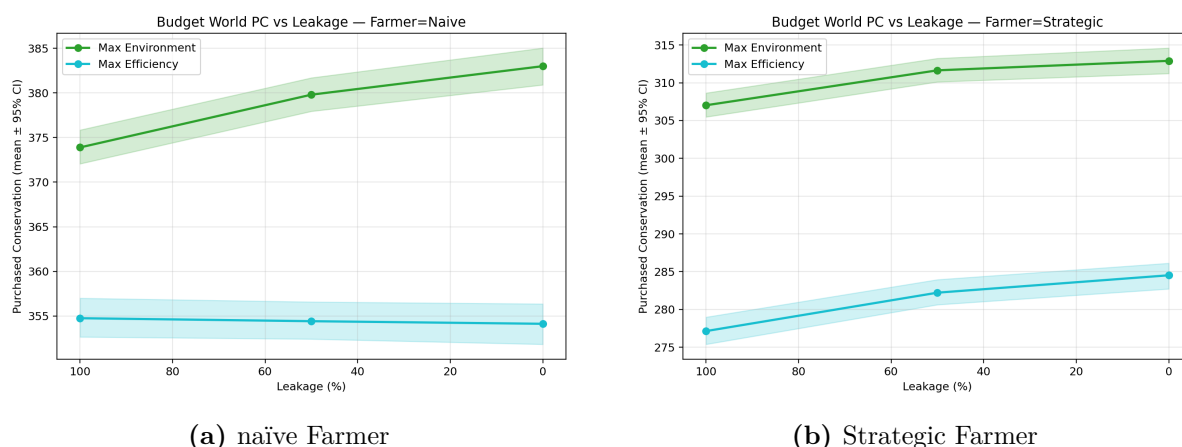


Figure A8.1: Budget World head-to-head Purchased Conservation (PC) versus leakage for naïve and Strategic Farmers. In the baseline simulations, MAX ENVIRONMENT strongly outperforms MAX EFFICIENCY; curves are nearly parallel across leakage, indicating that the selection effect is much larger than any residual sensitivity to L . Shaded bands: $\pm 95\%$ 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 Appendix A8.2 for the exact finite-instance discussion that accompanies these simulations.

A8.2.2 What can be shown exactly

Proposition A8.5 (Max Efficiency is not generally optimal). *Under equal sufficiently large budgets, there exist instances for which MAX EFFICIENCY (ratio-greedy e/a) yields strictly lower Purchased Conservation than MAX ENVIRONMENT (value-greedy e).*

Constructive example. We use the six-plot board from the introduction:

$$(e_i, a_i) \in \{(10, 9), (6, 10), (9, 7), (5, 3), (3, 2), (2, 4)\},$$

with equal budgets $B_g^0 = B_f^0 = 35$ and naïve Farmers (BAU ordering by descending a : (6, 10), (10, 9), (9, 7), (5, 3), (2, 4), (3, 2)).

Under MAX EFFICIENCY, Greens select by descending ratio. In round 1, Farmers take (6, 10); Greens take (5, 3). In round 2, Farmers take (10, 9); Greens take (3, 2). In round 3, Farmers take (9, 7); Greens take (2, 4). Purchased Conservation: PC = 5 + 3 + 2 = 10.

Under MAX ENVIRONMENT, Greens select by descending e . In round 1, Farmers take (6, 10); Greens take (10, 9). In round 2, Farmers take (9, 7); Greens take (5, 3). In round 3, Farmers take (2, 4); Greens take (3, 2). Purchased Conservation: PC = 10 + 5 + 3 = 18.

Thus MAX ENVIRONMENT strictly outperforms MAX EFFICIENCY on this instance. \square

Proposition A8.6 (No universal finite-instance dominance in the equal-budget procurement game). *Even with naïve Farmers, equal budgets, full observability, and full leakage, neither*

MAX ENVIRONMENT nor MAX EFFICIENCY weakly dominates the other on Purchased Conservation over all finite instances.

The next two constructions are stated in the broader equal-budget sequential procurement game. They are used only to show that universal finite-instance dominance fails under equal budgets; they are not presented as counterexamples inside the parity-calibrated benchmark $B_f^0 = B_g^0 = \frac{1}{2} \sum_i a_i$ used in Appendix A7.

Proof. Proposition A8.5 gives one direction. For the reverse direction, consider four plots

$$(e_i, a_i) \in \{(7, 9), (10, 6), (4, 11), (3, 1)\},$$

equal budgets $B_g^0 = B_f^0 = 13$, full leakage $L = 1$, and naïve Farmers. Farmers move first and purchase (4, 11). Under MAX ENVIRONMENT, Greens then purchase (10, 6), leaving $B_g = 7$. Farmers next purchase (3, 1), after which Greens cannot afford (7, 9). Hence

$$\text{PC}(\text{MAX ENVIRONMENT}) = 10.$$

Under MAX EFFICIENCY, Greens first purchase (3, 1), leaving $B_g = 12$. After buying (4, 11) the Farmer has only 2 units of budget left and cannot purchase either remaining plot. Greens then purchase (10, 6) on the next turn, so

$$\text{PC}(\text{MAX EFFICIENCY}) = 3 + 10 = 13.$$

Thus MAX EFFICIENCY also wins on some finite instances. □

Remark A8.7 (Interpretation). The exact finite game therefore does not support a blanket dominance theorem. The question is instead whether the paper’s benchmark simulation pattern reflects a structural regularity of the baseline data-generating process. The next subsection answers affirmatively.

A8.3 Asymptotic dominance: proof overviews

The counterexamples above establish that universal dominance fails. Theorem 4.1, stated in Section 4.1, proves that under the paper’s baseline DGP the knapsack reversal is nevertheless a genuine asymptotic result. At $\rho = 0.3$, the extension developed in Section A8.6 establishes the probability/frontier side and closes the affordability step by combining a validated ODE budget bound with a stopping-time bootstrap. This subsection provides the proof roadmap for the baseline and $\rho = 0.3$ cases. The full baseline proof is given in Section A8.4, the supporting input lemmas are collected in Section A8.5, and the correlated extension and its validated closure are given in Section A8.6.

Proof overview for Theorem 4.1. Under independence, the fluid-limit trajectory has an exact closed-form solution: the Farmer frontier and Green cutoff move in lockstep as $u(s) = v(s) = 1 - \sqrt{1 - 2s}$. This symmetry implies that Farmer’s per-turn ecological capture equals Green’s per-turn agricultural spend (Lemma A8.8), which together with the MX threshold identity yields the closed-form gap $\mu_E \phi - S_G(\tau) \approx 0.216$ as an exact asymptotic limit. Full details are in Section A8.4.

Proof overview for Theorem 4.2. The baseline architecture carries over with two new ingredients. First, a latent-coordinate survivor law and one-step drift calculation identify the doubly truncated bivariate-normal pool and Green’s actual per-turn value path. In those coordinates, Green’s per-turn ecological value converges to the deterministic function $v(s) = e(c(s))$, where $c(s)$ is the current Green cutoff on the latent ecological scale. Second, the cost side is closed by a validated ODE budget lower bound combined with a stopping-time affordability bootstrap, rather than by a static host-depletion inequality. The prefix wedge *increases* from +0.219 to +0.309 under positive correlation (the high-ratio corridor becomes more ecologically mediocre). The theorem then closes by comparing the fluid-limit value integral $\int_m^{\kappa_m^{\text{low}}} e(c(s)) ds$ to the MX threshold $\Theta_{0.3}$, with a validated margin of approximately 0.198. Full details are in Section A8.6.

Scope and limitations. The theorems are proved for the specific baseline DGP (uniform marginals on $[0.1, 10]$) at two correlation values. Simulations suggest that the normalised gap is positive across the entire range $\rho \in [0, 0.995]$, with a hump-shaped profile peaking near $\rho \approx 0.7$ (Figure 11), but only $\rho = 0$ and $\rho = 0.3$ carry formal proofs. The result is asymptotic; on very small grids (2×2), tight-budget counterexamples can dominate. The ecological selection premium (ESP) defined below provides a scalar diagnostic for when the theorem’s conditions are likely to hold.

A8.4 Baseline asymptotic dominance theorem

This subsection proves the baseline theorem for the Budget World object under independence:

$$(E_i, A_i)_{i=1}^N \text{ iid with marginals } U[0.1, 10], \quad E \perp A,$$

with parity budgets, Farmer moves first, full leakage, Purchased Conservation as the outcome, ME = affordable max- E , MX = affordable max- E/A .

The proof has a single idea: under $\rho = 0$ with identical marginals, the fluid-limit dynamics are symmetric in the two coordinates. This symmetry — unavailable at $\rho > 0$ — collapses the entire fluid-limit machinery to a one-line identity for the asymptotic gap.

A8.4.1 Fluid-limit trajectory under independence

In tail-probability coordinates $u(s) := 1 - F_U(x(s))$ and $v(s) := 1 - F_U(c(s))$ — where F_U denotes the CDF of $U[0.1, 10]$, so that u and v are the fractions of items above the current Farmer and Green frontiers — the fluid-limit ODE derived from the dynamic survivor law (Lemma A8.49) takes the form

$$\dot{u} = \frac{1}{1-v}, \quad \dot{v} = \frac{1}{1-u}, \quad \dot{b}_F = -a(x(s)), \quad \dot{b}_G = -g_A(x(s), c(s)),$$

with $u(0) = v(0) = 0$ and $b_F(0) = b_G(0) = \mu_A/2 = 2.525$. Under independence, the conditional-expectation kernel $g_A(x, c)$ reduces to the A -marginal conditional mean $\mathbb{E}[A \mid A \leq x]$, because $E \perp A$ makes Green’s ecological cutoff c irrelevant to the A -distribution

of the survivor pool.

The symmetry. Because the frontier drifts depend only on the complementary coordinate (\dot{u} on v and vice versa) and the initial conditions are identical, the solution satisfies $u(s) = v(s)$ for all s . The common trajectory solves $\dot{u} = 1/(1 - u)$, giving the explicit formula

$$(5) \quad u(s) = v(s) = 1 - \sqrt{1 - 2s}, \quad s \in [0, 1/2].$$

Substituting (5) into the budget ODEs and integrating gives closed-form cumulative spend paths

$$(6) \quad S_F(s) = 0.1 s + 3.3(1 - (1 - 2s)^{3/2}),$$

$$(7) \quad S_G(s) = 0.1 s + 1.65(1 - (1 - 2s)^{3/2}).$$

The ME Farmer-exhaustion share τ is the unique solution of $S_F(\tau) = 2.525$; interval evaluation gives $0.304718 < \tau < 0.304719$.

Complete allocation. After Farmer exhaustion, Green shops alone against a frozen frontier. A direct integration of the post-Farmer budget ODE (Appendix A8.5) yields the identity

$$(8) \quad \kappa = 1 - \tau,$$

so every item is purchased in the fluid limit. In the finite game, at most $\lceil 20/\ell \rceil = 200$ items remain unclaimed, giving $K_N^{\text{ME}}/N \rightarrow 1 - \tau$.

A8.4.2 The E–A symmetry and the closed-form gap

The key observation is that under $\rho = 0$ and identical marginals, each side’s per-turn capture on the other side’s dimension is distributionally identical.

Lemma A8.8 (E–A symmetry of per-turn captures). *Fix $s \in [0, \tau]$. Conditional on the fluid-limit state at share s , Farmer’s per-turn ecological value and Green’s per-turn agricultural*

cost have the same conditional expectation:

$$\mathbb{E}[E(f_{[sN]}^N)] = \mathbb{E}[A(g_{[sN]}^N)] + o(1) = \frac{1}{2}(\ell + x(s)),$$

where f_t^N and g_t^N denote Farmer's and Green's t -th picks. Consequently, Farmer's normalised cumulative ecological capture through share τ equals Green's normalised cumulative agricultural spend:

$$(9) \quad \frac{1}{N} \sum_{t=1}^{\lfloor \tau N \rfloor} E(f_t^N) \xrightarrow{p} \int_0^\tau \frac{\ell + x(s)}{2} ds = S_G(\tau).$$

Proof. By Lemma A8.49, immediately before Farmer's turn at share s , the survivor pool is iid from the doubly truncated host $\nu_{x(s^-), c(s^-)}$. Under $\rho = 0$, this host factors as a product of independent upper-truncated marginals: the A -coordinate is $U[\ell, x(s^-)]$ and the E -coordinate is $U[\ell, c(s^-)]$. Farmer picks $\arg \max A$; conditional on this pick's A -value being the truncation point $x(s^-) + o(1)$ in the fluid limit, its E -value is drawn from $U[\ell, c(s^-)]$ because $E \perp A$ in the host. Therefore

$$\mathbb{E}[E(f_{[sN]}^N)] = \frac{\ell + c(s)}{2} + o(1).$$

By the symmetric argument, Green picks $\arg \max E$ from the same factored host, and its A -value is drawn from $U[\ell, x(s^-)]$, giving $\mathbb{E}[A(g_{[sN]}^N)] = (\ell + x(s))/2 + o(1)$. The symmetry $x(s) = c(s)$ (equation (5) translated back to latent coordinates, using identical marginals) gives the displayed equality.

For (9), integrate the per-turn identity over $[0, \tau]$ and note that the right-hand side is

$$\int_0^\tau \mathbb{E}[A(g_{[sN]}^N)] ds = S_G(\tau) \text{ by definition of } S_G. \quad \square$$

Theorem A8.9 (Baseline asymptotic dominance theorem, exact form). *Under the baseline theorem object,*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}[\text{PC}_N^{\text{ME}} - \text{PC}_N^{\text{MX}}] = \mu_E \phi - S_G(\tau),$$

where ϕ is the MX Farmer-exhaustion share, τ is the ME Farmer-exhaustion share, and S_G is the Green agricultural spend path (7). Numerical evaluation gives

$$\mu_E \phi - S_G(\tau) = 1.493871 - 1.277736 = 0.216135 > 0.$$

Proof. We establish the identity by computing both sides asymptotically.

ME side. Total ecological value in the market satisfies $N^{-1} \sum_{i=1}^N E_i \xrightarrow{p} \mu_E$ by the law of large numbers. Lemma A8.13 gives exact play on the prefix $[0, m]$, Lemma A8.14 gives exact Green max- E play on every compact subinterval of $[m, \kappa)$, and Proposition A8.16 closes the endpoint by showing that only $O(1)$ labels remain unclaimed after the two budgets stop. Consequently, up to an $o_p(N)$ set of leftover labels, ME collects exactly the items not collected by Farmer, and hence

$$\frac{1}{N} \text{PC}_N^{\text{ME}} = \frac{1}{N} \sum_{i=1}^N E_i - \frac{1}{N} \sum_{t=1}^{K_N^F(\text{ME})} E(f_t^N) + o_p(1).$$

Proposition A8.15 gives $K_N^F(\text{ME})/N \xrightarrow{p} \tau$. Since per-item ecological values are bounded, Lemma A8.19 permits replacing the random upper limit by the deterministic one:

$$\frac{1}{N} \left| \sum_{t=1}^{K_N^F(\text{ME})} E(f_t^N) - \sum_{t=1}^{\lfloor \tau N \rfloor} E(f_t^N) \right| \xrightarrow{p} 0.$$

Lemma A8.8 then gives convergence of the second term to $S_G(\tau)$. Hence

$N^{-1} \text{PC}_N^{\text{ME}} \xrightarrow{p} \mu_E - S_G(\tau)$. Bounded convergence (per-item E -values lie in $[0.1, 10]$) upgrades this to convergence of expectations.

MX side. By Proposition A8.27, $K_N^F(\text{MX})/N \xrightarrow{p} \phi$. Under independence, MX's prefix ratio block and post-Farmer residual pool partition $\{A \leq q_F\}$ up to $o_p(N)$ items, where $q_F = Q_A(1 - \phi)$. By the LLN,

$$\frac{1}{N} \text{PC}_N^{\text{MX}} \xrightarrow{p} \mathbb{E}[E \cdot \mathbf{1}\{A \leq q_F\}] = \mu_E \cdot \mathbb{P}(A \leq q_F) = \mu_E(1 - \phi),$$

where the second equality uses $E \perp A$. Bounded convergence again gives convergence of expectations.

Combination. Subtracting,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}[\text{PC}_N^{\text{ME}} - \text{PC}_N^{\text{MX}}] = (\mu_E - S_G(\tau)) - \mu_E(1 - \phi) = \mu_E \phi - S_G(\tau).$$

Numerical evaluation. $\phi = (10 - \sqrt{50.005})/9.9 = 0.295816\dots$, so

$\mu_E \phi = 5.05 \times 0.295816 = 1.493871$. The bracket $0.304718 < \tau < 0.304719$ from $S_F(\tau) = 2.525$

gives, by the monotonicity of S_G ,

$$1.277734 < S_G(\tau) < 1.277738,$$

hence $0.216133 < \mu_E \phi - S_G(\tau) < 0.216137$, certified to five decimal places. \square

Remark A8.10 (Economic reading of the closed-form gap). The identity $\mu_E \phi - S_G(\tau)$ has a clean economic interpretation. MX concedes to Farmer all items with $A > q_F$ — exactly a ϕ -fraction of the market, each with average ecological value μ_E under independence. So MX's cession to Farmer is worth $\mu_E \phi$ in ecological terms. ME instead races Farmer, with both agents shopping from the same symmetric pool: by Lemma A8.8, the ecological value Farmer captures against ME equals the agricultural cost Green incurs against Farmer, which is $S_G(\tau)$. The gap is the difference: ME makes Farmer capture less ecological value because ME is selecting on the same dimension Farmer is indirectly depleting. The key mechanism is the symmetric fluid-limit dynamics under $\rho = 0$, which does not survive to $\rho > 0$, where Farmer's max- A pick is no longer E -uniform in the survivor pool.

Remark A8.11 (Why the general fluid-limit machinery is still needed). The identity of Theorem A8.9 relies essentially on the symmetry $x(s) = c(s)$, which fails under positive correlation: Farmer's max- A pick is correlated with the high- E region of the survivor pool, so Farmer's per-turn E capture exceeds Green's per-turn A spend, and no closed-form identity is available. The general fluid-limit architecture deployed in Appendix A8.6 — which tracks Green's actual per-turn ecological value through the ODE system (x, c, b_F, b_G) and closes the theorem via validated numerical integration — is the unified proof method. The baseline symmetry identity is a specialisation that happens to be available at $\rho = 0$.

Remark A8.12 (Certificate archival). A certification script (`baseline_closed_form_certificate.py`) using 80-digit interval arithmetic via `mpmath` is archived on Zenodo (DOI: 10.5281/zenodo.19598799). The certified asymptotic gap is $[0.216134, 0.216136]$.

A8.5 Supporting inputs S1–S4 for the baseline asymptotic dominance theorem

Frozen object, constants, and dependency map

We keep the notation from Appendix A8.4. Write

$$\ell := 0.1, \quad \nu := 10, \quad \mu_A := \mathbb{E}[A] = 5.05, \quad B_G^0 = B_F^0 = \frac{1}{2} \sum_{i=1}^N A_i.$$

The ratio score for MX is

$$R := E/A.$$

For any random variable Z with continuous distribution function F_Z , write

$$Q_Z(p) := \inf\{z : F_Z(z) \geq p\}, \quad p \in (0, 1),$$

for its population quantile. The Farmer exhaustion benchmark share $\phi = \phi_F^*$ is defined by

$$\int_{1-\phi}^1 Q_A(v) dv = \frac{\mu_A}{2}.$$

Under the baseline uniform law,

$$\phi \approx 0.295816, \quad q_F := Q_A(1 - \phi) \approx 7.07142.$$

The baseline asymptotic dominance theorem uses the following support items.

Item	Role in the baseline chain	Status here
S1	Empirical quantiles and thresholded LLNs for static portfolios and residual threshold functionals	standard results cited, forms stated precisely
S2	Replacing random split points (K_N^F, K_N^{ME}) by deterministic benchmarks $(\phi N, \kappa N)$	proved here
S3	Exactness of the initial ratio ordering for MX through Farmer exhaustion	proved here
S4	Fluid-limit theorem for the block-depletion chain producing $\alpha_\star(s)$ and $h_\star(s)$	standard theorem cited, conditions verified here

Lemma A8.13 (Early baseline affordability). *Let $a < \mu_A/(2\bar{u})$, where $\bar{u} = 10$ is the maximum item price. In the baseline theorem object, with probability tending to one both players can afford every plot through their first $\lfloor aN \rfloor$ scheduled moves. In particular, for the paper's prefix $m = 0.10$, Green and Farmer affordability are inactive on $[0, m]$.*

Proof. The initial normalized budgets satisfy

$$B_G^0/N, B_F^0/N \xrightarrow{p} \mu_A/2 = 2.525.$$

Through $\lfloor aN \rfloor$ own moves, either player can spend at most $\bar{u} aN + O(1)$, because every item price is at most \bar{u} . Since $a < \mu_A/(2\bar{u})$, choose $\gamma > 0$ with $\mu_A/2 - \bar{u}a > 2\gamma$. With probability tending to one, the residual normalized budget of each player is at least γ throughout the first $\lfloor aN \rfloor$ own moves. Hence the dollar budget is larger than 10 for all sufficiently large N on that event, and so every remaining plot is affordable. \square

Lemma A8.14 (Affordability is inactive on compact subintervals before κ). *Fix $\varepsilon \in (0, \kappa - m)$. Let*

$$x_\tau := Q_A(1 - u(\tau)) = 10 - 9.9u(\tau)$$

be the Farmer frontier at exhaustion, where $u(\cdot)$ is the active-phase trajectory in (5). Define

the fluid-limit Green-budget trajectory on $[m, \kappa]$ by

$$(10) \quad b_G(s) = \begin{cases} \mu_A/2 - S_G(s), & 0 \leq s \leq \tau, \\ \mu_A/2 - S_G(\tau) - (s - \tau) \frac{\ell + x_\tau}{2}, & \tau \leq s \leq \kappa. \end{cases}$$

Let

$$\beta_\varepsilon := \inf_{s \in [m, \kappa - \varepsilon]} b_G(s).$$

Then $\beta_\varepsilon > 0$, and for every fixed ε ,

$$\mathbb{P}(\text{Green plays exact arg max } Y \text{ throughout } [m, \kappa - \varepsilon]) \rightarrow 1 \quad \text{as } N \rightarrow \infty.$$

In particular, on that event the normalised chain converges uniformly to the fluid-limit trajectory on $[m, \kappa - \varepsilon]$.

Proof. The first line of (10) is the Farmer-active Green budget, obtained from the closed form (7). After Farmer exhaustion, the frontier x_τ is frozen. Under $\rho = 0$, the residual host factors, so Green's max- E pick has an A -coordinate drawn from the residual uniform marginal on $[\ell, x_\tau]$ up to an $o(1)$ endpoint error. The post-Farmer Green-only spending rate is therefore $(\ell + x_\tau)/2$, which gives the second line of (10). The two pieces agree at $s = \tau$. The identity $\kappa = 1 - \tau$ is equivalently the statement that this second line reaches zero at $s = \kappa$.

Thus b_G is continuous, strictly decreasing on each phase, and $b_G(\kappa) = 0$. Hence

$$\beta_\varepsilon = \inf_{s \in [m, \kappa - \varepsilon]} b_G(s) = b_G(\kappa - \varepsilon) > 0.$$

For the affordability claim we use a stopping-time bootstrap. Let

$$\sigma_N := \inf\{s \geq m : \text{some surviving plot has price} > b_G^N(s) N\}$$

denote the first affordability failure after the prefix, with $\sigma_N := +\infty$ if no failure occurs before κ . On $\{\sigma_N > s\}$, Green has played exact arg max Y throughout $[m, s]$. Lemma A8.49 then identifies the conditional survivor law at every Green move in that interval. On compact subintervals of $[m, \tau)$ the active-phase density-dependent chain has the drift used to derive (7);

on compact subintervals of (τ, κ) the Green-only chain has the frozen-frontier drift described above. The two deterministic budget paths concatenate continuously at τ , and the finite chain has bounded $O(1/N)$ increments in both phases.

Applying Kurtz's density-dependent LLN separately on the two phases and concatenating at the continuous phase boundary gives, on the stopped interval,

$$(11) \quad \sup_{m \leq s \leq (\kappa - \varepsilon) \wedge \sigma_N} \|Z_N(s) - Z(s)\| \xrightarrow{p} 0.$$

Suppose, for contradiction, that $\mathbb{P}(\sigma_N \leq \kappa - \varepsilon) \not\rightarrow 0$. On the event $\{\sigma_N \leq \kappa - \varepsilon\}$, by definition the maximum surviving plot price exceeds $b_G^N(\sigma_N)N$, and since every plot price is at most 10, this forces

$$b_G^N(\sigma_N) < 10/N.$$

But (11) forces $b_G^N(\sigma_N) \rightarrow b_G(\sigma_N \wedge (\kappa - \varepsilon)) \geq \beta_\varepsilon > 0$ in probability along any subsequence on which $\sigma_N \leq \kappa - \varepsilon$, contradicting $b_G^N(\sigma_N) \rightarrow 0$. Therefore $\mathbb{P}(\sigma_N > \kappa - \varepsilon) \rightarrow 1$, and on that event Green's rule is exact $\arg \max Y$ throughout $[m, \kappa - \varepsilon]$. Uniform convergence on the full interval then follows from (11) with σ_N replaced by $\kappa - \varepsilon$. \square

Proposition A8.15 (Baseline Farmer-exhaustion crossing time). *Let τ be the unique solution of $S_F(\tau) = \mu_A/2$ on $(0, 1/2)$, and let*

$$\tau_N := \frac{1}{N} \inf\{t : \text{Farmer cannot afford the current max-A item}\}$$

for the baseline ME path. If $K_N^F(\text{ME})$ denotes Farmer's total number of purchases on that path, then

$$\tau_N \xrightarrow{p} \tau, \quad \frac{K_N^F(\text{ME})}{N} \xrightarrow{p} \tau.$$

Proof. The random initial Farmer budget satisfies $B_F^0/N \rightarrow \mu_A/2$ in probability by the law of large numbers. Farmer-side affordability before τ follows by the same bootstrap argument applied to b_F^N . On any compact subinterval $[0, \tau - \varepsilon]$ with $\varepsilon > 0$, the fluid Farmer budget satisfies

$$\inf_{0 \leq s \leq \tau - \varepsilon} b_F(s) = b_F(\tau - \varepsilon) > 0.$$

Uniform convergence $b_F^N \rightarrow b_F$ on this subinterval, obtained from the active-phase density-dependent LLN used above with the Farmer budget coordinate retained, gives

$$b_F^N(s) \geq \frac{1}{2} b_F(\tau - \varepsilon) > 10/N$$

uniformly on $[0, \tau - \varepsilon]$ with probability tending to one. Since every item price is at most 10, Farmer can afford the current max- A item throughout $[0, \tau - \varepsilon]$ with probability tending to one.

The Farmer budget has a transversal zero at τ :

$$\dot{b}_F(\tau) = -(10 - 9.9 u(\tau)) \leq -0.1 < 0.$$

The continuous-mapping theorem for first-passage times of uniformly convergent monotone trajectories (Pollard [Pollard \(1984\)](#), Ch. 5) then gives $\tau_N \xrightarrow{p} \tau$. Letting $\varepsilon \downarrow 0$ in the compact-subinterval statement identifies the crossing time rather than only the pre-crossing affordability event.

It remains to connect this first max- A affordability failure to Farmer's total purchase count. At time $N\tau_N$, Farmer's residual budget is strictly smaller than the current maximum surviving agricultural price, and every price is at most 10. Hence the residual Farmer budget is below 10. Since every remaining plot costs at least $\ell = 0.1$, Farmer can make at most $\lceil 10/\ell \rceil = 100$ further purchases after this first failure. Thus, pathwise,

$$0 \leq K_N^F(\text{ME}) - N\tau_N \leq 100.$$

Dividing by N and using $\tau_N \xrightarrow{p} \tau$ proves $K_N^F(\text{ME})/N \xrightarrow{p} \tau$. □

Proposition A8.16 (Baseline leftover-count bound at terminal time). *Let T^* denote the terminal round at which neither player can afford any remaining plot. Under parity budgets $B_F^0 = B_G^0 = \frac{1}{2} \sum_i A_i$, the number of unclaimed items at T^* satisfies*

$$|U_{T^*}| \leq \left\lceil \frac{2\nu}{\ell} \right\rceil = 200.$$

Consequently $K_N^{\text{ME}}/N \xrightarrow{p} 1 - \tau$.

Proof. Parity budgets give the pathwise identity

$$B_F^{\text{res}}(t) + B_G^{\text{res}}(t) = \sum_{i \in U_t} A_i$$

at every round t , because the two initial budgets sum to the total agricultural cost $\sum_{i=1}^N A_i$ and every deletion reduces the corresponding residual by exactly the item's price. At the terminal round T^* , neither player can afford any remaining plot; in particular, $B_F^{\text{res}}(T^*) < \nu$ and $B_G^{\text{res}}(T^*) < \nu$. Therefore

$$\sum_{i \in U_{T^*}} A_i = B_F^{\text{res}}(T^*) + B_G^{\text{res}}(T^*) < 2\nu = 20.$$

Since each surviving plot has $A_i \geq \ell = 0.1$, $|U_{T^*}| \leq \lceil 2\nu/\ell \rceil = 200$ is immediate.

For the consequence, Proposition A8.15 gives $K_N^F(\text{ME})/N \xrightarrow{P} \tau$, and Lemma A8.14 shows that Green is unconstrained throughout the Farmer-active phase and the Green-only phase up to $\kappa - \varepsilon$, for every $\varepsilon > 0$. Since $\kappa = 1 - \tau > \tau$, with probability tending to one Farmer exhausts strictly before Green: the terminal round T^* coincides with Green's stopping time, and Farmer is already unable to afford any remaining plot at T^* . The leftover bound then reads

$$N - K_N^{\text{ME}} - K_N^F(\text{ME}) = |U_{T^*}| \leq 200$$

on that event. Dividing by N and combining with $K_N^F(\text{ME})/N \xrightarrow{P} \tau$ yields

$$K_N^{\text{ME}}/N \xrightarrow{P} 1 - \tau. \quad \square$$

Appendix A8.4 contains the direct outside-prefix comparison that closes the baseline theorem. This appendix collects the supporting inputs behind that proof.

S1. Empirical quantile and thresholded-LLN inputs

The baseline theorem repeatedly replaces empirical threshold portfolios by their population limits. The relevant classes are simple: one-dimensional quantile cutoffs, ratio half-spaces $\{R \geq r\}$, and residual slices of the form

$$\{A \leq q, t \leq R < r\}.$$

These are VC classes with bounded envelope under the compact support $[\ell, \nu]^2$. So standard Glivenko–Cantelli results apply to the indicator class itself and to bounded weights 1, A , and E .

Proposition A8.17 (Quantile and thresholded-LLN facts used in the baseline proof). *Let $(Z_i)_{i \geq 1}$ be iid with continuous cdf F_Z and quantile function Q_Z . Let $\hat{Q}_{Z,N}$ be the empirical quantile function. Then, for every compact $K \subset (0, 1)$,*

$$\sup_{p \in K} |\hat{Q}_{Z,N}(p) - Q_Z(p)| \xrightarrow{a.s.} 0.$$

Now let $(E_i, A_i)_{i \geq 1}$ be iid on $[\ell, \nu]^2$ with continuous law, and write $R_i = E_i/A_i$. We use only thresholds whose boundary sets have probability zero; under the continuous laws in this paper this holds for $A = q$, $R = r$, and finite intersections of such boundaries. Let g be any bounded measurable function on $[\ell, \nu]^2$. Then:

(i) for every compact interval $I \subset (0, \infty)$,

$$\sup_{r \in I} \left| \frac{1}{N} \sum_{i=1}^N g(E_i, A_i) \mathbf{1}\{R_i \geq r\} - \mathbb{E}[g(E, A) \mathbf{1}\{R \geq r\}] \right| \xrightarrow{a.s.} 0;$$

(ii) for every compact rectangle $K \subset [\ell, \nu] \times (0, \infty)^2$,

$$\sup_{(q,t,r) \in K} \left| \frac{1}{N} \sum_{i=1}^N g(E_i, A_i) \mathbf{1}\{A_i \leq q, t \leq R_i < r\} - \mathbb{E}[g(E, A) \mathbf{1}\{A \leq q, t \leq R < r\}] \right| \xrightarrow{a.s.} 0;$$

(iii) if r_s is the unique solution of $\mathbb{P}(R \geq r_s) = s$, then

$$\frac{1}{N} \sum_{i=1}^N g(E_i, A_i) \mathbf{1}\{R_i \geq \hat{r}_{s,N}\} \xrightarrow{a.s.} \mathbb{E}[g(E, A) \mathbf{1}\{R \geq r_s\}],$$

where $\hat{r}_{s,N}$ is the empirical top- s ratio cutoff.

Proof / reference. The empirical quantile statement is standard; see, for example, van der Vaart (? , Chapter 21) or Shorack and Wellner (? , Chapter 2). The indicator classes

$$\{\mathbf{1}\{R \geq r\} : r > 0\} \quad \text{and} \quad \{\mathbf{1}\{A \leq q, t \leq R < r\} : (q, t, r) \in K\}$$

are VC classes on $[\ell, \nu]^2$, so bounded weighted versions are Glivenko–Cantelli. The third claim follows by combining the first two parts with continuity of the ratio law and the convergence $\hat{r}_{s,N} \rightarrow r_s$. \square

Remark A8.18 (Where S1 is used). S1 is the justification behind all static portfolio functionals in the baseline chain. In particular it identifies

$$M_X(s) := \mathbb{E}[E\mathbf{1}\{R \geq r_s\}], \quad C_X(s) := \mathbb{E}[A\mathbf{1}\{R \geq r_s\}],$$

as the large- N value and cost of the initial top- sN ratio portfolio, and it also underpins the residual threshold maps t_x , $C_X^{\text{res}}(x; \phi)$, and $V_X^{\text{res}, E}(x; \phi)$.

S2. Benchmark-split reductions

The baseline theorem is written at deterministic split points such as $\lfloor \phi N \rfloor$ and $\lfloor \kappa N \rfloor$, while the underlying game has random stopping times. The normalization by N makes this harmless whenever split times converge and per-turn terms are bounded.

Lemma A8.19 (Deterministic replacement of convergent split times). *Let τ_N be an integer-valued random time with $\tau_N/N \rightarrow \tau \in [0, 1]$ in probability. Let $(Y_{N,t})_{1 \leq t \leq N}$ be random variables with $|Y_{N,t}| \leq M$ almost surely for all N, t . Then*

$$\frac{1}{N} \mathbb{E} \left| \sum_{t=1}^{\tau_N} Y_{N,t} - \sum_{t=1}^{\lfloor \tau N \rfloor} Y_{N,t} \right| \rightarrow 0.$$

The same conclusion holds for window sums: if $\sigma_N/N \rightarrow \sigma$ in probability, then

$$\frac{1}{N} \mathbb{E} \left| \sum_{t=\tau_N+1}^{\sigma_N} Y_{N,t} - \sum_{t=\lfloor \tau N \rfloor+1}^{\lfloor \sigma N \rfloor} Y_{N,t} \right| \rightarrow 0.$$

Proof. Pathwise,

$$\left| \sum_{t=1}^{\tau_N} Y_{N,t} - \sum_{t=1}^{\lfloor \tau N \rfloor} Y_{N,t} \right| \leq M |\tau_N - \lfloor \tau N \rfloor|.$$

Fix $\eta > 0$. On the event $|\tau_N/N - \tau| \leq \eta$, the right-hand side is at most $MN\eta + M$. On the

complement it is at most MN . Therefore

$$\frac{1}{N} \mathbb{E} \left| \sum_{t=1}^{\tau_N} Y_{N,t} - \sum_{t=1}^{\lfloor \tau_N \rfloor} Y_{N,t} \right| \leq M\eta + \frac{M}{N} + M\mathbb{P}(|\tau_N/N - \tau| > \eta).$$

Take $N \rightarrow \infty$ and then $\eta \downarrow 0$. The window-sum version follows by applying the same argument separately to the upper and lower endpoints. \square

Remark A8.20 (Where S2 is used). S2 is the mechanism for replacing actual random split times—such as the Farmer exhaustion time or the ME stopping time—by their deterministic limits inside normalized value or cost sums. This is not deep; it just needs to be stated instead of waved away.

S3. Safe-frontier exactness through Farmer exhaustion

Appendix A8.4 already established safe-frontier separation at $m = 0.10$. What the theorem note actually uses later is the stronger statement that the same structural exactness persists through the benchmark Farmer exhaustion share ϕ .

Population separation at the benchmark share Under baseline independence and uniform marginals, the ratio tail has the closed form

$$(12) \quad \mathbb{P}(R \geq r) = \frac{(\nu - \ell r)^2}{2r(\nu - \ell)^2}, \quad 1 \leq r \leq \nu/\ell.$$

Let r_s solve $\mathbb{P}(R \geq r_s) = s$. Then $s \mapsto r_s$ is decreasing. At the benchmark exhaustion share,

$$r_\phi \approx 1.66752, \quad \frac{\nu}{r_\phi} \approx 5.99752 < q_F \approx 7.07142.$$

Because $s \mapsto \nu/r_s$ is increasing and $s \mapsto Q_A(1 - s)$ is decreasing, the worst case on $[0, \phi]$ is precisely $s = \phi$. Hence

$$(13) \quad \frac{\nu}{r_s} < Q_A(1 - s) \quad \text{for every } s \in [0, \phi].$$

Equation (13) is the global safe-frontier inequality on the whole Farmer-active window.

Proposition A8.21 (Safe-frontier exactness on compact subintervals and at the endpoint).

Fix the baseline independent-uniform law and a deterministic $T < \phi$. Let $m_\phi(N) = \lfloor \phi N \rfloor$.

Then there exists an event $\Omega_N^{\text{sf}}(T)$ with $\mathbb{P}(\Omega_N^{\text{sf}}(T)) \rightarrow 1$ such that on $\Omega_N^{\text{sf}}(T)$:

- (i) the initial top- $m_\phi(N)$ ratio block is disjoint from the initial top- $m_\phi(N)$ Farmer A-block;
- (ii) for every $1 \leq t \leq \lfloor TN \rfloor$, the Farmer's first t purchases under the MX path all lie outside the initial top- $m_\phi(N)$ ratio block;
- (iii) for every $1 \leq t \leq \lfloor TN \rfloor$, the t -th MX purchase is exactly the t -th item in the static initial descending ratio order.

Consequently,

$$\frac{1}{N} \sum_{t=1}^{\lfloor TN \rfloor} E(g_t^{\text{MX}}) \xrightarrow{p} M_X(T), \quad \frac{1}{N} \sum_{t=1}^{\lfloor TN \rfloor} A(g_t^{\text{MX}}) \xrightarrow{p} C_X(T).$$

Moreover, the endpoint identities hold in the compact-to-endpoint sense:

$$\frac{1}{N} \sum_{t=1}^{m_\phi(N)} E(g_t^{\text{MX}}) \xrightarrow{p} M_X(\phi), \quad \frac{1}{N} \sum_{t=1}^{m_\phi(N)} A(g_t^{\text{MX}}) \xrightarrow{p} C_X(\phi).$$

The same conclusion holds for every fixed subwindow $[m, \phi] \subset [0, \phi]$.

Proof. Choose $\eta > 0$ small enough that

$$\frac{\nu}{r_\phi - 2\eta} < q_F - 2\eta.$$

Let $\Omega_N^{\text{sf}}(T)$ be the event that empirical A- and ratio-cutoffs are within η of their population counterparts and that both players are unconstrained through $\lfloor TN \rfloor$ turns. The empirical-cutoff part of $\Omega_N^{\text{sf}}(T)$ has probability tending to one by Proposition A8.17. Farmer affordability through this window follows from the definition of the no-Green descending-A crossing time σ_N and the crossing-time LLN $\sigma_N/N \rightarrow \phi$, since $T < \phi$. Green affordability through this window follows from Corollary A8.26, which applies on an event of probability tending to one so long as $T < \phi + \delta_{\text{aff}}$; this holds for any $T < \phi$. Therefore $\mathbb{P}(\Omega_N^{\text{sf}}(T)) \rightarrow 1$.

On $\Omega_N^{\text{sf}}(T)$, every item in the initial top- $m_\phi(N)$ ratio block has ratio at least $r_\phi - \eta$, hence

agricultural value at most

$$\frac{\nu}{r_\phi - \eta} < q_F - \eta.$$

Meanwhile every item in the initial top- $m_\phi(N)$ Farmer A -block has agricultural value at least $q_F - \eta$. Therefore the two initial blocks are disjoint, proving (i).

Now argue by induction on $t \leq \lfloor TN \rfloor$. At time $t = 1$, the Farmer chooses inside the initial top- $m_\phi(N)$ A -block and therefore outside the ratio block. Suppose the claim holds through round $t - 1$. Then Farmer has deleted only items outside the initial ratio block, while MX has deleted only items inside that block. So among the surviving labels, the highest remaining ratio item is simply the next item in the static initial ratio order. That proves (iii) at time t . Since the Farmer has still not touched the ratio block, his t -th purchase again lies outside it, proving (ii) and closing the induction.

The convergence of normalized static sums up to T follows from S1. To pass to ϕ , fix $T < \phi$. Since $E, A \in [\ell, \nu]$,

$$\frac{1}{N} \left| \sum_{t=\lfloor TN \rfloor+1}^{m_\phi(N)} E(g_t^{\text{MX}}) \right| \leq \nu(\phi - T) + o(1),$$

and the same deterministic bound holds with $A(g_t^{\text{MX}})$ in place of $E(g_t^{\text{MX}})$. The population tails satisfy

$$|M_X(\phi) - M_X(T)| + |C_X(\phi) - C_X(T)| \leq 2\nu(\phi - T).$$

Thus, after taking $N \rightarrow \infty$ for the fixed compact interval and then letting $T \uparrow \phi$, the endpoint identities follow. Applying the same argument to two endpoints gives the subwindow statement. \square

Remark A8.22 (Why S3 is stronger than the prefix lemma). The point is not just that the top tenth of the ratio portfolio is safe. The point is that the whole initial top- ϕN ratio aisle stays untouched by the Farmer on every compact Farmer-active subinterval, and therefore survives to the benchmark exhaustion split after normalization. That is what makes the post-Farmer residual pool D_ϕ a fixed static region rather than a path-dependent mess.

Remark A8.23 (Compact-to-endpoint upgrade). Proposition A8.21 deliberately proves pathwise exactness first on each fixed $T < \phi$. The endpoint is obtained by the deterministic

tail bound written in the proof:

$$\frac{1}{N} \left| \sum_{t=\lfloor TN \rfloor + 1}^{\lfloor \phi N \rfloor} Y_{N,t} \right| \leq 10(\phi - T) + o(1)$$

for any per-turn quantity $Y_{N,t} \in [0, 10]$, followed by $T \uparrow \phi$. The residual pool D_ϕ and the objects $M_X(\phi)$, $C_X(\phi)$ are therefore justified as compact-limit objects of the actual MX game paths, not by a separate unproved endpoint jump.

MX-side Green affordability. The safe-frontier argument that follows establishes *disjointness* of the static top- A block from the static top-ratio block. That alone does not imply MX actually purchases the top-ratio block: without an affordability guarantee, MX could be budget-constrained and forced to take a lower-ratio item that might lie in Farmer's top- A block. We close that gap with a distributional lemma and a stopping-time corollary.

Lemma A8.24 (MX-side Green affordability, baseline). *Under the baseline independent-uniform law, the residual cost function*

$$C_R(r) := \mathbb{E}[A \mathbf{1}\{R \geq r\}] = \frac{\nu^3/(6r^2) - \nu\ell^2/2 + r\ell^3/3}{d^2}, \quad r \in [1, \nu/\ell],$$

is continuous and strictly decreasing in r . At the benchmark exhaustion share $\phi \approx 0.29582$ and $r_\phi \approx 1.66752$, substitution gives $C_R(r_\phi) \approx 0.6110 < \mu_A/2 = 2.525$. Consequently there exist constants $\delta_{\text{aff}} > 0$ and $\varepsilon_0 > 0$ such that

$$C_R(r_{\phi+\delta}) \leq \frac{\mu_A}{2} - \varepsilon_0, \quad \text{for every } \delta \in [0, \delta_{\text{aff}}].$$

Proof. For $r \in [1, \nu/\ell]$ the event $\{R \geq r\}$ is the triangle $\{(a, e) \in [\ell, \nu]^2 : e \geq ra\}$, which forces $a \leq \nu/r$. Integrating the joint uniform density $1/d^2$,

$$C_R(r) = \frac{1}{d^2} \int_{\ell}^{\nu/r} a(\nu - ra) da = \frac{1}{d^2} \left[\frac{\nu a^2}{2} - \frac{r a^3}{3} \right]_{\ell}^{\nu/r},$$

which simplifies to the displayed closed form. Differentiating under the integral via Leibniz's rule (the boundary term vanishes because the integrand $a(\nu - ra)$ is zero at $a = \nu/r$),

$$C'_R(r) = -\frac{1}{d^2} \int_{\ell}^{\nu/r} a^2 da < 0,$$

so C_R is strictly decreasing in r on $[1, \nu/\ell]$. The conclusion then follows from continuity: since C_R is continuous and $s \mapsto r_s$ is continuous in s , and since $C_R(r_\phi) \approx 0.6110 < \mu_A/2 = 2.525$ with margin roughly 1.9, there exists $\delta_{\text{aff}} > 0$ and $\varepsilon_0 > 0$ such that $C_R(r_{\phi+\delta}) \leq \mu_A/2 - \varepsilon_0$ for all $\delta \in [0, \delta_{\text{aff}}]$. \square

Lemma A8.25 (Baseline safe-frontier buffer). *Under the baseline independent-uniform law, the strict safe-frontier inequality*

$$\frac{\nu}{r_\phi} < q_F$$

holds, with $\nu/r_\phi \approx 5.998 < q_F \approx 7.071$. Consequently there exist constants $\delta_{\text{sf}} > 0$ and $\eta_{\text{sf}} > 0$ such that

$$\frac{\nu}{r_{\phi+\delta_{\text{sf}}} - 2\eta_{\text{sf}}} < Q_A(1 - \phi - \delta_{\text{sf}}) - 2\eta_{\text{sf}}.$$

Let $\Omega_N^{\text{sf}} := \Omega_N^{\text{sf}}(\delta_{\text{sf}}, \eta_{\text{sf}})$ be the event that the empirical ratio- and A -cutoffs at share $\phi + \delta_{\text{sf}}$ are within η_{sf} of their population counterparts. Then $\Pr(\Omega_N^{\text{sf}}) \rightarrow 1$, and on Ω_N^{sf} the initial top- $(\phi + \delta_{\text{sf}})N$ ratio block is disjoint from the first $\lfloor (\phi + \delta_{\text{sf}})N \rfloor$ items of the descending A -order.

Proof. The baseline ratio law has explicit density, giving $r_\phi \approx 1.66752$ via $\Pr(R \geq r_\phi) = \phi$, and $q_F = Q_A(1 - \phi) \approx 7.07142$; direct substitution yields $\nu/r_\phi \approx 5.998 < 7.071 \approx q_F$. By continuity of $s \mapsto r_s$ and $s \mapsto Q_A(1 - s)$, the buffer constants $\delta_{\text{sf}}, \eta_{\text{sf}} > 0$ satisfying the displayed strict inequality exist. Proposition A8.17 then gives $\Pr(\Omega_N^{\text{sf}}) \rightarrow 1$. Disjointness on Ω_N^{sf} is the same corner argument as in the proof of Proposition A8.27: every item in the initial top- $(\phi + \delta_{\text{sf}})N$ ratio block has $A \leq \nu/\hat{r} \leq \nu/(r_{\phi+\delta_{\text{sf}}} - \eta_{\text{sf}}) < Q_A(1 - \phi - \delta_{\text{sf}}) - \eta_{\text{sf}}$, whereas every item in the first $\lfloor (\phi + \delta_{\text{sf}})N \rfloor$ descending A -items has $A \geq \hat{q} \geq Q_A(1 - \phi - \delta_{\text{sf}}) - \eta_{\text{sf}}$. \square

Corollary A8.26 (Green is unconstrained through the MX safe-frontier window, baseline).

Fix δ_{aff} and ε_0 as in Lemma A8.24, and $\delta_{\text{sf}}, \eta_{\text{sf}}$ as in Lemma A8.25. Set $\delta_0 := \min(\delta_{\text{sf}}, \delta_{\text{aff}})$ and $T_N := \lfloor (\phi + \delta_0)N \rfloor$. Writing $B_{G,N}^{\text{res}}(t)$ for Green's absolute residual budget before round t (in dollars), on an event Ω_N^{afford} of probability tending to one,

$$B_{G,N}^{\text{res}}(t) \geq \frac{\varepsilon_0 N}{2} > \nu \quad \text{for every } t \leq T_N \text{ and every } N \text{ sufficiently large.}$$

Consequently Green plays exact arg max R over the entire surviving host throughout rounds $1, \dots, T_N$ (not just the affordable subset).

Proof. Define the Green affordability failure time

$$\tau_N^G := \inf \left\{ t \geq 1 : \max_{i \in S_t} A_i > B_{G,N}^{\text{res}}(t) \right\}, \quad \tau_N^G := \infty \text{ otherwise,}$$

where S_t denotes the set of plots surviving before round t . On $\{\tau_N^G > t\} \cap \Omega_N^{\text{sf}}$, with Ω_N^{sf} the event from Lemma A8.25, Green's first t picks are exactly the first t items of the initial descending ratio order: Green is unconstrained (by definition of τ_N^G) so plays exact $\arg \max R$ over the whole surviving host, and disjointness from Farmer's top- A block (Lemma A8.25) means no item in the ratio block has been removed by Farmer. Their cumulative agricultural cost satisfies

$$\sum_{k=1}^t A(g_k^{\text{MX}}) \leq \sum_{i: R_i \geq \hat{r}_{\phi+\delta_0, N}} A_i.$$

By Proposition A8.17,

$$\frac{1}{N} \sum_{i: R_i \geq \hat{r}_{\phi+\delta_0, N}} A_i \xrightarrow{p} C_R(r_{\phi+\delta_0}) \leq \frac{\mu_A}{2} - \varepsilon_0.$$

With $B_G^0/N \xrightarrow{p} \mu_A/2$, Green's absolute residual budget on $\{\tau_N^G > t\} \cap \Omega_N^{\text{sf}}$ satisfies

$B_{G,N}^{\text{res}}(t) \geq \varepsilon_0 N/2$ for every $t \leq T_N$, with probability tending to one. For N large enough that $\varepsilon_0 N/2 > \nu$, the maximum surviving plot price at any $t \leq T_N$ is strictly less than $B_{G,N}^{\text{res}}(t)$.

This contradicts $\{\tau_N^G \leq T_N\}$. Therefore $\tau_N^G > T_N$ on

$$\Omega_N^{\text{afford}} := \Omega_N^{\text{sf}} \cap \left\{ B_{G,N}^{\text{res}}(T_N) \geq \varepsilon_0 N/2 \right\},$$

which has probability tending to one. □

Proposition A8.27 (MX endpoint exhaustion at the benchmark share). *Under the baseline law,*

$$\frac{K_N^F(\text{MX})}{N} \xrightarrow{p} \phi.$$

Proof. Write $A_{(1)} \geq \dots \geq A_{(N)}$ for the initial descending A -order. Define the frontier time at which the no-Green descending- A path first leaves less than 10 units of budget:

$$\sigma_N := \min \left\{ m \geq 0 : B_F^0 - \sum_{j=1}^m A_{(j)} < 10 \right\},$$

with the convention $\sum_{j=1}^0 A_{(j)} = 0$.

Let

$$S_A(s) := \int_0^s Q_A(1-u) du.$$

By Proposition A8.17, for each fixed $s \in (0, 1)$,

$$\frac{1}{N} \sum_{j=1}^{\lfloor sN \rfloor} A_{(j)} \xrightarrow{p} S_A(s),$$

and by the ordinary law of large numbers,

$$\frac{B_F^0}{N} = \frac{1}{2N} \sum_{i=1}^N A_i \xrightarrow{p} \frac{\mu_A}{2}.$$

Since $S_A(\phi) = \mu_A/2$, S_A is continuous and strictly increasing, and $10/N \rightarrow 0$, the standard crossing-time argument gives

$$\frac{\sigma_N}{N} \xrightarrow{p} \phi.$$

Next use the safe-frontier margin. At the benchmark share we have the strict inequality

$$\frac{\nu}{r_\phi} < q_F, \quad q_F := Q_A(1-\phi).$$

By continuity of $s \mapsto r_s$ and $s \mapsto Q_A(1-s)$, choose a fixed $\delta > 0$ such that

$$\frac{\nu}{r_{\phi+\delta}} < Q_A(1-\phi-\delta).$$

Let $\hat{r}_{\phi+\delta, N}$ be the empirical top- $(\phi+\delta)$ ratio cutoff, and let $\hat{q}_{\phi+\delta, N}$ be the empirical top- $(\phi+\delta)$ A -cutoff. Choose $\eta > 0$ so small that

$$\frac{\nu}{r_{\phi+\delta} - 2\eta} < Q_A(1-\phi-\delta) - 2\eta.$$

Define the event

$$\Omega_N^{\text{sf}} := \{|\hat{r}_{\phi+\delta, N} - r_{\phi+\delta}| \leq \eta, |\hat{q}_{\phi+\delta, N} - Q_A(1-\phi-\delta)| \leq \eta\}.$$

By Proposition A8.17, $\mathbb{P}(\Omega_N^{\text{sf}}) \rightarrow 1$.

On Ω_N^{sf} , every item in the initial top- $\lfloor(\phi + \delta)N\rfloor$ ratio block has

$$A \leq \frac{\nu}{\hat{r}_{\phi+\delta, N}} \leq \frac{\nu}{r_{\phi+\delta} - \eta} < Q_A(1 - \phi - \delta) - \eta,$$

whereas each of the first $\lfloor(\phi + \delta)N\rfloor$ items in the initial descending A -order has

$$A \geq \hat{q}_{\phi+\delta, N} \geq Q_A(1 - \phi - \delta) - \eta.$$

So the initial top- $(\phi + \delta)N$ ratio block is disjoint from the first $\lfloor(\phi + \delta)N\rfloor$ items in the descending A -order.

Now work on the event

$$\Omega_N^{\text{sf}} \cap \Omega_N^{\text{afford}} \cap \{\sigma_N \leq (\phi + \delta)N\},$$

where Ω_N^{afford} is the affordability event from Corollary A8.26 (we shrink δ if necessary so that $\delta \leq \delta_{\text{aff}}$; the two conditions are both open, so a common δ exists). We claim that, up to round σ_N , the actual Farmer path under MX coincides with the no-Green descending- A path.

Indeed, before round $t \leq \sigma_N$, the Farmer's remaining budget is at least 10, so the highest surviving A -item is affordable. By induction, Farmer's first $t - 1$ purchases are

$A_{(1)}, \dots, A_{(t-1)}$. By Corollary A8.26, on Ω_N^{afford} Green plays exact $\arg \max R$ over the full affordable host throughout rounds $1, \dots, \lfloor(\phi + \delta)N\rfloor$ without being budget-constrained, so every earlier Green purchase lies in the initial top- $\lfloor(\phi + \delta)N\rfloor$ ratio block. Hence none of $A_{(t)}, \dots, A_{(\lfloor(\phi+\delta)N\rfloor)}$ has been removed, so the current highest surviving A -item is still $A_{(t)}$.

Therefore Farmer's t -th purchase is $A_{(t)}$. This proves the claim.

Consequently, on $\Omega_N^{\text{sf}} \cap \Omega_N^{\text{afford}} \cap \{\sigma_N \leq (\phi + \delta)N\}$, Farmer makes exactly σ_N purchases before his remaining budget first falls below 10. After that moment he can make at most 100 further purchases, because every remaining plot costs at least 0.1. Thus

$$0 \leq K_N^F(\text{MX}) - \sigma_N \leq 100 \quad \text{on } \Omega_N^{\text{sf}} \cap \Omega_N^{\text{afford}} \cap \{\sigma_N \leq (\phi + \delta)N\}.$$

Since $\sigma_N/N \rightarrow \phi$ and

$$\mathbb{P}(\Omega_N^{\text{sf}} \cap \Omega_N^{\text{afford}} \cap \{\sigma_N \leq (\phi + \delta)N\}) \rightarrow 1,$$

it follows that

$$\frac{K_N^F(\text{MX})}{N} \xrightarrow{p} \phi.$$

□

S4. Density-dependent LLN for the block-depletion chain

The acceleration envelope and the mirror E -frontier both rest on the same one-dimensional density-dependent chain. The theorem itself is standard; what matters here is verifying that our chain satisfies its hypotheses.

Theorem A8.28 (Specialized density-dependent LLN; quoted). *Let $(Z_t^N)_{t \geq 0}$ be a Markov chain with values in a compact interval $K \subset \mathbb{R}$, and define the scaled-time process*

$$z_N(s) = Z_{\lfloor sN \rfloor}^N. \text{ Suppose that on } [0, T]:$$

(i) *one-step increments satisfy $|Z_{t+1}^N - Z_t^N| \leq C/N$ almost surely;*

(ii) *the conditional drift has the form*

$$\mathbb{E}[Z_{t+1}^N - Z_t^N \mid \mathcal{F}_t] = \frac{1}{N} f(Z_t^N, t/N) + o(N^{-1})$$

uniformly on compact subsets of $K \times [0, T]$;

(iii) *f is locally Lipschitz on the region visited by the limit path;*

(iv) *$Z_0^N \rightarrow z_0$ in probability.*

Then the ODE

$$\dot{z}(s) = f(z(s), s), \quad z(0) = z_0,$$

has a unique solution on $[0, T]$, and

$$\sup_{0 \leq s \leq T} |z_N(s) - z(s)| \xrightarrow{p} 0.$$

Reference. This is a standard special case of the fluid-limit theorems of Kurtz [Kurtz \(1970\)](#); [Ethier and Kurtz \(1986\)](#) or Wormald's differential-equation method [Wormald \(1999\)](#). The theorem is quoted rather than reproved. □

Proposition A8.29 (Verification for the baseline block-depletion chain). *Fix $T < 1/2$ and $\alpha \in (0, 1)$. Under the natural “before Green turn t ” state timing, let $X_t^{N,\alpha}$ be the unnormalised count of surviving labels in the block under consideration and set*

$$x_t^{N,\alpha} := \frac{1}{N} X_t^{N,\alpha}.$$

The count recursion is

$$X_{t+1}^{N,\alpha} = X_t^{N,\alpha} - \mathbf{1}\{X_t^{N,\alpha} > 0\} - \xi_t^{N,\alpha},$$

where, conditional on the current history,

$$\mathbb{E}[\xi_t^{N,\alpha} \mid \mathcal{F}_t] = \frac{X_t^{N,\alpha} - \mathbf{1}\{X_t^{N,\alpha} > 0\}}{N - 2t} + O\left(\frac{1}{N}\right).$$

Equivalently, on every compact positive region, the normalized chain satisfies

$$\mathbb{E}[x_{t+1}^{N,\alpha} - x_t^{N,\alpha} \mid \mathcal{F}_t] = \frac{1}{N} \left(-1 - \frac{x_t^{N,\alpha}}{1 - 2t/N} \right) + o(N^{-1}).$$

Consequently $x_N^\alpha(s) := x_{[sN]}^{N,\alpha}$ obeys Theorem A8.28 with limiting drift

$$f(x, s) = -1 - \frac{x}{1 - 2s}$$

and limit

$$(14) \quad x'_\alpha(s) = -1 - \frac{x_\alpha(s)}{1 - 2s}, \quad x_\alpha(0) = \alpha,$$

whenever the limit stays strictly positive.

Proof. For the normalized chain,

$$|x_{t+1}^{N,\alpha} - x_t^{N,\alpha}| \leq 2/N,$$

so condition (i) in Theorem A8.28 is immediate. The normalized state space is compact because $0 \leq x_t^{N,\alpha} \leq 1$. On any region where the count stays above a positive floor, the

indicator equals one and the conditional drift is

$$-\frac{1}{N} - \frac{1}{N} \frac{x_t^{N,\alpha}}{1 - 2t/N} + O(N^{-2}),$$

which yields the displayed drift function. The map $f(x, s) = -1 - x/(1 - 2s)$ is locally Lipschitz on $[c, 1] \times [0, T]$ for every $c > 0$, and the normalized initial condition satisfies $x_0^{N,\alpha} \rightarrow \alpha$. So Theorem A8.28 applies on any compact interval on which the limit path stays strictly positive. \square

Corollary A8.30 (Closed form, frontier, and acceleration profile). *The ODE (14) has explicit solution*

$$x_\alpha(s) = (1 - 2s) + (\alpha - 1)\sqrt{1 - 2s}.$$

The critical boundary where this solution hits zero is

$$\alpha_\star(s) = 1 - \sqrt{1 - 2s},$$

and the corresponding acceleration profile is

$$h_\star(s) := \alpha_\star(s) - s = 1 - \sqrt{1 - 2s} - s.$$

In particular,

$$h_\star(0.10) \approx 0.00557, \quad h_\star(\phi) \approx 0.06514, \quad \int_{0.10}^{\phi} h_\star(s) ds \approx 0.00554.$$

Proof. Direct solution of the linear ODE, followed by the boundary condition

$$x_{\alpha_\star(s)}(s) = 0. \quad \square$$

Remark A8.31 (A-side versus mirror E-side). The A-side acceleration chain and the mirror E-side chain are not the same economic object, but after the state-timing normalization above they generate the same fluid drift. That is why the same α_\star and h_\star appear in both the baseline frontier convergence and the $\rho = 0.3$ frontier envelope.

A8.6 Positive-correlation extension at $\rho = 0.3$: frontier envelope and validated closure

This subsection studies the same Budget World theorem object as Appendix A8.4, but under the Gaussian-copula law with latent correlation $\rho = 0.3$. We use short forms only to keep formulas readable: ME denotes *Max Environment*, MX denotes *Max Efficiency*, the paper's ecological and agricultural values are written E_i and A_i , and the price of plot i is A_i .

Purchased Conservation is the Green team's direct purchase total,

$$\text{PC}_N^H := \sum_{t=1}^{K_N^H} E(g_t^H), \quad H \in \{\text{ME}, \text{MX}\},$$

where g_t^H is the t -th Green purchase under rule H , and K_N^H is the total number of Green purchases under that rule.

Object	Direct $\rho = 0.3$ extension object
Values and price	$(E_i, A_i)_{i=1}^N$ are iid from the Gaussian-copula law with latent correlation $\rho = 0.3$, uniform marginals on $[0.1, 10]$, and price of plot i equal to A_i .
Budgets	Budget parity: $B_G^0 = B_F^0 = \frac{1}{2} \sum_{i=1}^N A_i$.
Leakage	Full leakage: $L = 1$. Farmer budget is not burned when Greens buy a Farmer-BAU plot.
Farmer rule	Naïve Farmer, moving first, buys the highest affordable remaining A -item.
Green rules	ME: highest affordable E (equivalently, highest Y on the latent scale). MX: highest affordable ratio $R := E/A$.
Outcome	Purchased Conservation PC. This is the Budget World analogue of the paper's PSE channel; DLE is not part of the theorem object.
Fixed split point	Prefix share $m = 0.10$.

On the latent Gaussian scale, write

$$A = 0.1 + 9.9\Phi(X), \quad E = 0.1 + 9.9\Phi(Y), \quad Y = \rho X + \sqrt{1 - \rho^2}\varepsilon, \quad \rho = 0.3,$$

with X, ε independent standard normals. Since E is a strictly increasing transform of Y , the ME rule is equivalently $\arg \max Y$.

The point of this subsection is twofold. First, it establishes the probabilistic foundation for the correlated Farmer-active frontier and value argument. Second, it closes the remaining affordability step by combining a validated ODE budget lower bound with a stopping-time bootstrap. Together, these prove the full direct $\rho = 0.3$ extension.

A8.6.1 A diagnostic decomposition and the $\rho = 0.3$ proof target

Fix $m \in (0, 1/2)$ and let $m_N := \lfloor mN \rfloor$. Because the ME and MX paths can stop at different Green purchase counts, the literal algebraic decomposition must be written with the common stopping window

$$K_N^\wedge := \min\{K_N^{\text{ME}}, K_N^{\text{MX}}\}.$$

For any realization,

$$\begin{aligned} \text{PC}_N^{\text{ME}} - \text{PC}_N^{\text{MX}} &= \sum_{t=1}^{m_N} \left(E(g_t^{\text{ME}}) - E(g_t^{\text{MX}}) \right) + \sum_{t=m_N+1}^{K_N^\wedge} \left(E(g_t^{\text{ME}}) - E(g_t^{\text{MX}}) \right) \\ (15) \quad &+ \mathbf{1}\{K_N^{\text{ME}} > K_N^{\text{MX}}\} \sum_{t=K_N^{\text{MX}}+1}^{K_N^{\text{ME}}} E(g_t^{\text{ME}}) - \mathbf{1}\{K_N^{\text{MX}} > K_N^{\text{ME}}\} \sum_{t=K_N^{\text{ME}}+1}^{K_N^{\text{MX}}} E(g_t^{\text{MX}}). \end{aligned}$$

This is the exact identity. The shorter three-term display sometimes used in diagnostic tables corresponds to the special case $K_N^{\text{MX}} \geq K_N^{\text{ME}}$ and to a common-window indexing convention; it is not used as a formal proof step below.

For $\rho = 0.3$, it is cleaner to close the extension by direct post-prefix comparison rather than by separately optimizing the common-window and terminal-tail terms in (15). Define

$$V_{E,0.3}^{\text{out}} := \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left[\sum_{t=m_N+1}^{K_N^{\text{ME}}} E(g_t^{\text{ME}}) \right],$$

the post-prefix ME ecological value. Using Proposition A8.36, the $M_X(0.10, 0.3)$ term cancels algebraically in the ME-vs-MX difference, and the rigorous closure uses the certified ceiling from Proposition A8.64,

$$(16) \quad \Theta_{0.3} := \mathbb{E}[E \mathbf{1}\{A \leq q_F\}] - 0.901 \leq 2.36413.$$

Thus it suffices to prove $V_{E,0.3}^{\text{out}} > 2.36413$.

The remaining subsections track the Farmer-active ecological value through the global fluid-limit value path and close the affordability contribution by a validated ODE budget bound and stopping-time bootstrap.

Remark A8.32 (Proof architecture). Equation (15) provides useful diagnostic bookkeeping, but the formal proof does not close through separate bounding of its middle and terminal-tail components. Instead, the theorem is proved by tracking Green's actual ecological value path in the fluid limit, which captures the full game dynamics without worst-case deletion bounds on the value side. The affordability side uses the validated ODE budget bound and stopping-time bootstrap in Proposition A8.60; the value side uses the global fluid-limit integral. See Section A8.6.6 below.

A8.6.2 Safe-frontier geometry and the exact MX side

Let $R := E/A$, and for $s \in (0, 1)$ let $r_s(0.3)$ solve

$$\mathbb{P}(R \geq r_s(0.3)) = s.$$

At $\rho = 0.3$, the accompanying interval-arithmetic certificates give

$$(17) \quad r_{0.10}(0.3) \in [3.62406450, 3.62406459], \quad r_\phi(0.3) \in [1.50517830, 1.50517845].$$

We display midpoint values in prose only for readability; all inequalities below use the conservative interval endpoints.

MX-side Green affordability at $\rho = 0.3$. Exactly the same gap that appears in the baseline MX proof reappears here: safe-frontier disjointness does not by itself imply that Green can afford the top-ratio block. We close the gap with a copula analogue of Lemma A8.24.

Lemma A8.33 (MX-side Green affordability, $\rho = 0.3$). *Under the Gaussian-copula law with $\rho = 0.3$ and uniform marginals on $[\ell, \nu]$, the residual cost function*

$$C_R^{(0.3)}(r) := \mathbb{E}_{0.3}[A \mathbf{1}\{R \geq r\}]$$

is continuous and strictly decreasing in r . Outward-rounded interval quadrature on the ratio half-space $\{R \geq r\}$ under the Gaussian-copula joint density, using the same one-sided Abramowitz–Stegun Φ -bounds as the scripts cited in Appendix A8.6.7, establishes

$$(18) \quad C_R^{(0.3)}(r_{\phi}(0.3)) \in [0.679, 0.685] < \frac{\mu_A}{2} = 2.525.$$

Consequently there exist constants $\delta_{\text{aff}} > 0$ and $\varepsilon_0 \geq 1.83$ such that

$$C_R^{(0.3)}(r_{\phi+\delta}(0.3)) \leq \frac{\mu_A}{2} - \varepsilon_0, \quad \text{for every } \delta \in [0, \delta_{\text{aff}}].$$

Proof. Continuity and strict monotonicity of $C_R^{(0.3)}$ follow from the absolute continuity of the Gaussian-copula law with respect to Lebesgue measure on $[\ell, \nu]^2$. The interval (18) is produced by latent-coordinate interval quadrature of $\int(\ell + d\Phi(x))\Phi((\rho x - \Phi^{-1}(r\Phi(x) + \ell(r-1)/d))/\sigma)\varphi(x)dx$ over $\{x : r\Phi(x) + \ell(r-1)/d < 1\}$, using the one-sided Abramowitz–Stegun Φ -bounds, verified inverse-normal brackets, and outward-rounded IEEE arithmetic described in (CI1)–(CI2) of Appendix A8.6.7. The quadrature is evaluated at the certified endpoints $r \in [1.50517830, 1.50517845]$ from (17), using r_{lo} for the upper bound (larger integration region, larger integrand) and r_{hi} for the lower bound. The displayed bracket is conservative: refining the quadrature mesh tightens the enclosure interval, but the margin $\varepsilon_0 > 1.83$ is already more than two orders of magnitude larger than any required safety threshold.¹⁰ The margin constants δ_{aff} and ε_0 then follow from continuity of $C_R^{(0.3)}$ and r_s in s . □

Lemma A8.34 ($\rho = 0.3$ safe-frontier buffer). *Under the Gaussian-copula law with $\rho = 0.3$,*

¹⁰The displayed interval is extracted from the machine-readable certificate report under the key `mx_affordability_rho03`.

the strict safe-frontier inequality

$$\frac{10}{r_\phi(0.3)} < q_F$$

holds. Consequently there exist constants $\delta_{\text{sf}} > 0$ and $\eta_{\text{sf}} > 0$ such that

$$\frac{10}{r_{\phi+\delta_{\text{sf}}}(0.3) - 2\eta_{\text{sf}}} < Q_A(1 - \phi - \delta_{\text{sf}}) - 2\eta_{\text{sf}},$$

and the corresponding empirical-cutoff event $\Omega_N^{\text{sf}} := \Omega_N^{\text{sf}}(\delta_{\text{sf}}, \eta_{\text{sf}})$ satisfies $\Pr(\Omega_N^{\text{sf}}) \rightarrow 1$. On Ω_N^{sf} , the initial top- $(\phi + \delta_{\text{sf}})N$ ratio block is disjoint from the first $\lfloor (\phi + \delta_{\text{sf}})N \rfloor$ items of the descending A -order.

Proof. Using the interval enclosures in (17),

$$q_F - \frac{10}{\underline{r}_\phi} \geq 7.0714215 - \frac{10}{1.50517830} > 0.42769 > 0,$$

where \underline{r}_ϕ is the certified lower endpoint of the ratio-cut interval at share ϕ . The ratio law is continuous under the Gaussian-copula model and the marginal A -law is uniform, so buffer constants $\delta_{\text{sf}}, \eta_{\text{sf}} > 0$ exist by continuity. Proposition A8.17 gives $\Pr(\Omega_N^{\text{sf}}) \rightarrow 1$, and the disjointness follows from the same corner argument as in Lemma A8.25. \square

Corollary A8.35 (Green is unconstrained through the MX safe-frontier window, $\rho = 0.3$).

Fix δ_{aff} and ε_0 as in Lemma A8.33, and $\delta_{\text{sf}}, \eta_{\text{sf}}$ as in Lemma A8.34. Set $\delta_0 := \min(\delta_{\text{sf}}, \delta_{\text{aff}})$.

Writing $B_{G,N}^{\text{res}}(t)$ for Green's absolute residual budget before round t (in dollars), on an event Ω_N^{afford} of probability tending to one,

$$B_{G,N}^{\text{res}}(t) \geq \frac{\varepsilon_0 N}{2} > \nu \quad \text{for every } t \leq \lfloor (\phi + \delta_0)N \rfloor \text{ and every } N \text{ sufficiently large.}$$

Consequently Green plays exact $\arg \max R$ over the entire surviving host throughout rounds $1, \dots, \lfloor (\phi + \delta_0)N \rfloor$ (not just the affordable subset).

Proof. Identical to the proof of Corollary A8.26, with the baseline C_R replaced by $C_R^{(0.3)}$, the baseline Ω_N^{sf} replaced by the $\rho = 0.3$ event from Lemma A8.34, and the baseline affordability lemma replaced by Lemma A8.33. \square

Proposition A8.36 (Safe-frontier exactness through ϕ). Under the Gaussian-copula law with

$\rho = 0.3$, let $r_s(0.3)$ be defined by $\Pr(R \geq r_s(0.3)) = s$, and let $q_F := Q_A(1 - \phi)$. Then the strict safe-frontier inequality

$$\frac{10}{r_\phi(0.3)} < q_F$$

holds. Consequently, on an event whose probability tends to one:

- (i) MX's prefix purchases are the initial static top-ratio portfolio through the share $m = 0.10$;
- (ii) through Farmer exhaustion, MX's Green purchases agree with the initial descending ratio order up to $o_p(N)$ endpoint labels, and $K_N^F(\text{MX})/N \rightarrow \phi$ in probability;
- (iii) after Farmer exhaustion, the empirical residual pool converges, for every bounded threshold-continuity test function, to the static region

$$D_{\phi,0.3} := \{(e, a) : a \leq q_F, e/a < r_\phi(0.3)\}.$$

In particular, normalized MX value and cost sums may be computed from the static top-ratio block through ϕ and the residual region $D_{\phi,0.3}$.

Proof. The safe-frontier condition at share s is

$$\frac{e_+}{r_s(0.3)} < Q_A(1 - s), \quad e_+ = 10.$$

Using the interval enclosures in (17),

$$Q_A(0.9) - \frac{10}{\underline{r}_{0.10}} \geq 9.01 - \frac{10}{3.62406450} > 6.25066 > 0,$$

and, at the benchmark Farmer-exhaustion share,

$$q_F - \frac{10}{\underline{r}_\phi} \geq 7.0714215 - \frac{10}{1.50517830} > 0.42769 > 0,$$

where $\underline{r}_{0.10}$ and \underline{r}_ϕ denote the certified lower endpoints of the two ratio-cut intervals. The ratio law is continuous under the Gaussian-copula model, and the marginal A -law is uniform.

Hence, by continuity of $s \mapsto r_s(0.3)$ and $s \mapsto Q_A(1 - s)$, choose $\delta > 0$ and $\eta > 0$ such that

$$(19) \quad \frac{10}{r_{\phi+\delta}(0.3) - 2\eta} < Q_A(1 - \phi - \delta) - 2\eta.$$

Let $\hat{r}_{\phi+\delta, N}$ be the empirical top- $(\phi + \delta)N$ ratio cutoff and let $\hat{q}_{\phi+\delta, N}$ be the empirical top- $(\phi + \delta)N$ A -cutoff. Let Ω_N^{sf} be the event on which

$$|\hat{r}_{\phi+\delta, N} - r_{\phi+\delta}(0.3)| \leq \eta, \quad |\hat{q}_{\phi+\delta, N} - Q_A(1 - \phi - \delta)| \leq \eta,$$

and the no-Green descending- A budget crossing time

$$\sigma_N := \min \left\{ k \geq 0 : B_F^0 - \sum_{j=1}^k A_{(j)} < 10 \right\}$$

satisfies $\sigma_N \leq (\phi + \delta)N$. Here $A_{(1)} \geq \dots \geq A_{(N)}$ is the initial descending A -order. The empirical quantile laws and the ordinary budget LLN give $\Pr(\Omega_N^{\text{sf}}) \rightarrow 1$, because $N^{-1}\sigma_N \rightarrow \phi$.

On Ω_N^{sf} , every item in the initial top- $\lfloor (\phi + \delta)N \rfloor$ ratio block has

$$A \leq \frac{10}{\hat{r}_{\phi+\delta, N}} \leq \frac{10}{r_{\phi+\delta}(0.3) - \eta} < Q_A(1 - \phi - \delta) - \eta,$$

while every item among $A_{(1)}, \dots, A_{(\lfloor (\phi+\delta)N \rfloor)}$ has

$$A \geq \hat{q}_{\phi+\delta, N} \geq Q_A(1 - \phi - \delta) - \eta.$$

Thus the initial top- $(\phi + \delta)N$ ratio block and the first $(\phi + \delta)N$ descending- A items are disjoint.

We now prove the path statement by induction. Without loss, shrink δ if necessary so that $\delta \leq \delta_{\text{aff}}$; the safe-frontier and affordability conditions are both open, so a common $\delta > 0$ exists, and by Corollary A8.35, on $\Omega_N^{\text{sf}} \cap \Omega_N^{\text{afford}}$ Green plays exact $\arg \max R$ over the entire surviving host throughout rounds $1, \dots, \lfloor (\phi + \delta)N \rfloor$. Suppose the claim holds through round $t - 1 \leq \sigma_N - 1$: Farmer has bought $A_{(1)}, \dots, A_{(t-1)}$, and MX has bought the first $t - 1$ items in the static descending ratio order. Since these two static blocks are disjoint on Ω_N^{sf} , none of $A_{(t)}, \dots, A_{(\lfloor (\phi+\delta)N \rfloor)}$ has been removed by MX. Also, by definition of σ_N , Farmer's residual budget before round $t \leq \sigma_N$ is at least 10, so the current highest surviving A -item is affordable

and is exactly $A_{(t)}$. Farmer buys it. Farmer has not touched the ratio block, and by the corollary Green is unconstrained through round $\lfloor(\phi + \delta)N\rfloor$, so the next MX purchase is the next item in the initial static ratio order. This proves the induction through σ_N .

At σ_N , Farmer's remaining budget is below 10, and each further Farmer purchase costs at least $\ell = 0.1$. Therefore

$$0 \leq K_N^F(\text{MX}) - \sigma_N \leq \lceil 10/\ell \rceil = 100$$

on $\Omega_N^{\text{sf}} \cap \Omega_N^{\text{afford}}$. Since $\sigma_N/N \rightarrow \phi$, this proves $K_N^F(\text{MX})/N \rightarrow \phi$. The same induction with any fixed $T < \phi$ gives exact static ratio ordering up to $\lfloor TN \rfloor$; Lemma A8.19 and the preceding convergence of $K_N^F(\text{MX})/N$ replace T by ϕ in normalized value and cost sums.

Finally, the strict inequality $10/r_\phi(0.3) < q_F$ implies $\{R \geq r_\phi(0.3)\} \subset \{A < q_F\}$. The Farmer-active Green block is therefore the static top-ratio block inside $\{A \leq q_F\}$, while Farmer's active block is the empirical top- A block $\{A > q_F\}$, up to quantile boundary labels. The labels left for MX after Farmer exhaustion are accordingly the empirical version of $D_{\phi,0.3} = \{A \leq q_F, R < r_\phi(0.3)\}$, up to $o_p(N)$ boundary labels. Thresholded LLNs then give the stated static residual-pool limit for all bounded test functions whose discontinuity set has population measure zero. \square

Distribution-free prefix lower bound for ME. The ME prefix lower envelope \underline{M}_E does not depend on the copula: it relies only on the ecological marginal law and a deterministic alternating order-statistic bound. We state the two inputs explicitly here because an earlier draft invoked a convenient but *incorrect* pathwise half-top inequality. The valid replacement is an alternating-rank bound combined with a uniform-spacing lemma.

Lemma A8.37 (Alternating order-statistic bound on the ME prefix). *Let*

$E_{(1)} \geq E_{(2)} \geq \dots \geq E_{(N)}$ *be the ecological order statistics of the realized draw, and let* g_t^{ME} *denote Green's* t -*th purchase under Max Environment. On the event that, through Green's first* $\lfloor mN \rfloor$ *turns, Green can afford every surviving plot and therefore selects the surviving plot with maximal* E , *the pathwise lower bound*

$$\sum_{t=1}^{\lfloor mN \rfloor} E(g_t^{\text{ME}}) \geq \sum_{t=1}^{\lfloor mN \rfloor} E_{(2t)}$$

holds on every realization.

Proof. Immediately before Green's t -th move, at most $2t - 1$ plots have been removed from the host: the Farmer has made at most t moves (including any Farmer move completed in round t before Green acts), and Green has made $t - 1$ earlier moves. Hence among the top- $2t$ ecological plots $\{E_{(1)}, \dots, E_{(2t)}\}$, at least one survives in the host. By the event definition, that surviving top- $2t$ plot is affordable, and because Green selects the surviving plot with maximal E (over the entire surviving host, not just an affordable subset), Green's t -th choice has ecological value at least $E_{(2t)}$. Summing over $1 \leq t \leq \lfloor mN \rfloor$ gives the claim. \square

Lemma A8.38 (Adjacent order-statistic spacing under iid uniform marginals). *Let E_1, \dots, E_N be iid $U[\ell, \nu]$. The joint law of (E_i, A_i) within each plot may be arbitrary, provided the draws (E_i, A_i) are iid across plots and the E -marginal is uniform. Then, for every fixed $m \in (0, 1/2)$,*

$$\frac{1}{N} \sum_{t=1}^{\lfloor mN \rfloor} (E_{(2t-1)} - E_{(2t)}) \xrightarrow{p} 0,$$

and consequently

$$\frac{1}{N} \sum_{t=1}^{\lfloor mN \rfloor} E_{(2t)} = \frac{1}{2N} \sum_{j=1}^{2\lfloor mN \rfloor} E_{(j)} + o_p(1) \xrightarrow{p} \frac{1}{2} \int_{1-2m}^1 Q_E(u) du.$$

Proof. Write $E_i = \ell + dU_i$ with $U_i \sim U[0, 1]$ iid, and let $U_{(1)} < \dots < U_{(N)}$ be the ascending uniform order statistics. Then

$$E_{(2t-1)} - E_{(2t)} = d(U_{(N-2t+2)} - U_{(N-2t+1)}).$$

Each uniform spacing has expectation $1/(N + 1)$. The random variable

$$Z_N := \frac{1}{N} \sum_{t=1}^{\lfloor mN \rfloor} (E_{(2t-1)} - E_{(2t)})$$

is nonnegative with

$$\mathbb{E}Z_N = \frac{d}{N} \cdot \frac{\lfloor mN \rfloor}{N + 1} = O(N^{-1}),$$

so Markov's inequality gives $Z_N \xrightarrow{p} 0$. The algebraic identity

$$\sum_{j=1}^{2\lfloor mN \rfloor} E_{(j)} = 2 \sum_{t=1}^{\lfloor mN \rfloor} E_{(2t)} + \sum_{t=1}^{\lfloor mN \rfloor} (E_{(2t-1)} - E_{(2t)})$$

holds pathwise for every N . Replacing $2\lfloor mN \rfloor$ by $\lfloor 2mN \rfloor$ changes the normalized sum by at most ν/N (they differ by at most one order statistic, each bounded by ν), so the quantile LLN $N^{-1} \sum_{j=1}^{\lfloor 2mN \rfloor} E_{(j)} \xrightarrow{p} \int_{1-2m}^1 Q_E(u) du$ from Proposition A8.17 applies equally to $N^{-1} \sum_{j=1}^{2\lfloor mN \rfloor} E_{(j)}$. Combining yields the displayed limit. \square

Remark A8.39 (Why these replace the previous half-top inequality). An earlier draft asserted that ME's first mN purchases dominate one half of the empirical top- $2mN$ ecological sum *pathwise*. That inequality is false in general: if $E_{(1)} \gg E_{(2)}$, Farmer can delete $E_{(1)}$ and leave Green with $E_{(2)}$, whereas $(E_{(1)} + E_{(2)})/2$ can be arbitrarily larger than $E_{(2)}$. The correct deterministic inequality is Lemma A8.37, in which Green's t -th purchase is lower-bounded by $E_{(2t)}$ rather than by an average of adjacent pairs. The half-top value is recovered *asymptotically* by the spacing lemma, which shows that the difference between the alternating sum and one half of the top- $2mN$ sum vanishes after normalization by N .

Proposition A8.40 (Prefix wedge at $\rho = 0.3$). *Let*

$$\underline{M}_E(m) := \frac{1}{2} \int_{1-2m}^1 Q_E(u) du, \quad M_X(m, 0.3) := \mathbb{E}[E \mathbf{1}\{R \geq r_m(0.3)\}],$$

and let the prefix PC gap be

$$G_{N,m}^{\text{PC}} := \sum_{t=1}^{\lfloor mN \rfloor} (E(g_t^{\text{ME}}) - E(g_t^{\text{MX}})).$$

Then

$$\underline{M}_E(0.10) = 0.901, \quad M_X(0.10, 0.3) \in [0.59221770, 0.59221775],$$

and the prefix PC gap satisfies the asymptotic lower bound

$$(20) \quad \liminf_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}[G_{N,0.10}^{\text{PC}}] \geq 0.901 - 0.59221775 > 0.3087822.$$

Proof. ME lower envelope. Through Green's first $mN = 0.10N$ turns, the early-affordability lemma (Lemma A8.13; the argument uses only the marginal price bound $A \leq \nu$ and the

normalized-budget LLN, so it applies identically under the Gaussian-copula law) gives, with probability tending to one, that Green can afford every surviving plot:

$0.10\nu = 1 < \mu_A/2 = 2.525$ leaves an order- N budget margin. On that event Green selects the surviving plot with maximal E through round mN , so Lemma A8.37 yields the pathwise inequality

$$\sum_{t=1}^{\lfloor mN \rfloor} E(g_t^{\text{ME}}) \geq \sum_{t=1}^{\lfloor mN \rfloor} E_{(2t)}.$$

The marginal law of E remains $U[\ell, \nu]$ regardless of ρ and the draws are iid across plots, so Lemma A8.38 converts the right-hand side into the asymptotic lower envelope

$$\frac{1}{N} \sum_{t=1}^{\lfloor mN \rfloor} E(g_t^{\text{ME}}) \geq \underline{M}_E(m) + o_p(1), \quad \underline{M}_E(m) = \frac{1}{2} \int_{1-2m}^1 Q_E(u) du.$$

Direct integration at $m = 0.10$ with $Q_E(u) = 0.1 + 9.9u$ gives

$$\underline{M}_E(0.10) = \frac{1}{2} \int_{0.8}^1 (0.1 + 9.9u) du = \frac{1}{2}(5.05 - 3.248) = 0.901.$$

MX prefix value. By Proposition A8.36, Green's first $\lfloor mN \rfloor$ picks under MX coincide with the empirical top- m ratio portfolio, up to $o_p(N)$ boundary labels. The ratio half-space family $\{R \geq r\}$ is a one-parameter VC-class with bounded E -envelope on the compact support $[\ell, \nu]^2$, so the thresholded LLN of Proposition A8.17 applies under the Gaussian-copula joint law.

Hence

$$\frac{1}{N} \sum_{t=1}^{\lfloor mN \rfloor} E(g_t^{\text{MX}}) \xrightarrow{p} M_X(m, 0.3) = \mathbb{E}[E \mathbf{1}\{R \geq r_m(0.3)\}].$$

Certified interval quadrature of the copula density on the ratio half-space yields

$$M_X(0.10, 0.3) \in [0.59221770, 0.59221775].^{11}$$

Conclusion. Combining the ME lower envelope with the conservative upper endpoint of the MX prefix,

$$\frac{1}{N} G_{N,0.10}^{\text{PC}} \geq 0.901 - 0.59221775 + o_p(1).$$

Uniform boundedness of per-item E -values in $[\ell, \nu]$ delivers uniform integrability of the normalized prefix sums, so the in-probability inequality upgrades to the displayed expectation lower bound (20). □

¹¹Extracted from the machine-readable certificate report under the key `mx_prefix_value_rho03`.

Remark A8.41 (Diagnostic corroborative numerics, not used in the theorem chain). The following rounded interval numerics are recorded for diagnostic cross-checks only; they do not enter the closure of Theorem 4.2, which uses the single ceiling $\Theta_{0.3} \leq 2.36413$ from Proposition A8.64 rather than the decomposition into prefix gap plus outside-MX value.

Under $\rho = 0.3$, quadrature intervals consistent with the archived (CI1)–(CI5) primitives give

$$M_X(\phi, 0.3) - M_X(0.10, 0.3) \in [1.23010750, 1.23010768],$$

$$V_{X,0.3}^{\text{res},E}(x_{\max}; \phi) \in [1.44279265, 1.44279283], \quad x_{\max} = 1 - 2\phi \approx 0.408368,$$

adding to a diagnostic upper bound $V_{X,0.3}^{\text{out}} \lesssim 2.67290051$. Because these intervals are corroborative rather than part of the theorem chain, they are not archived as separate keys in the certificate reports.

A8.6.3 The ME Farmer-active frontier and value path

Under positive correlation, Farmer deletions are no longer E -neutral, so the baseline mirror- E frontier argument does not apply directly. The proof instead works on the latent Gaussian scale, conditioning on the surviving X -profile and the realized past Green picks, and then applying the monotone-argmax lemma conditionally.

Lemma A8.42 (Monotone argmax for upper-truncated normals). *Fix a common upper truncation level $c \in \mathbb{R} \cup \{+\infty\}$, a variance $\sigma^2 > 0$, and means*

$$\mu_1 \leq \mu_2 \leq \cdots \leq \mu_M.$$

Let $Y_i^{(c)}$ be independent normals $N(\mu_i, \sigma^2)$, truncated above at c . Then the probabilities

$$\mathbb{P}(Y_i^{(c)} = \max_j Y_j^{(c)})$$

are nondecreasing in μ_i . Consequently, for any subset consisting of the k largest means, the argmax falls in that subset with probability at least k/M .

Proof. For $c \in \mathbb{R} \cup \{+\infty\}$, let f_μ^c denote the density of the normal law $N(\mu, \sigma^2)$ truncated to

$(-\infty, c]$. Thus, with the convention $\Phi(+\infty) = 1$,

$$f_{\mu}^c(y) = \frac{1}{\sigma \Phi((c - \mu)/\sigma)} \varphi\left(\frac{y - \mu}{\sigma}\right) \mathbf{1}\{y \leq c\}.$$

If $\mu' < \mu$, then for $y \leq c$,

$$\frac{f_{\mu}^c(y)}{f_{\mu'}^c(y)} = \frac{\Phi((c - \mu')/\sigma)}{\Phi((c - \mu)/\sigma)} \exp\left(\frac{\mu - \mu'}{\sigma^2} y - \frac{\mu^2 - (\mu')^2}{2\sigma^2}\right).$$

Its log-derivative in y is

$$\frac{d}{dy} \log\left(\frac{f_{\mu}^c(y)}{f_{\mu'}^c(y)}\right) = \frac{\mu - \mu'}{\sigma^2} \geq 0.$$

So the family $\{f_{\mu}^c : \mu \in \mathbb{R}\}$ has monotone likelihood ratio in y , even after truncation to the common set $(-\infty, c]$.

Now write

$$p_i := \mathbb{P}(Y_i^{(c)} = \max_{1 \leq j \leq M} Y_j^{(c)}),$$

and let F_i denote the cdf corresponding to $f_i := f_{\mu_i}^c$. Because the truncated normal laws are continuous, ties occur with probability zero, and therefore

$$p_i = \int_{-\infty}^c f_i(y) \prod_{\ell \neq i} F_{\ell}(y) dy.$$

Fix $i < j$. Then

$$p_j - p_i = \int_{-\infty}^c (f_j(y)F_i(y) - f_i(y)F_j(y)) \prod_{\ell \neq i, j} F_{\ell}(y) dy.$$

Set

$$r(y) := \frac{f_j(y)}{f_i(y)}.$$

By the monotone-likelihood-ratio property, r is nondecreasing on $(-\infty, c]$. Hence for every $y \leq c$,

$$F_j(y) = \int_{-\infty}^y f_j(u) du = \int_{-\infty}^y r(u) f_i(u) du \leq r(y) \int_{-\infty}^y f_i(u) du = \frac{f_j(y)}{f_i(y)} F_i(y).$$

Therefore

$$f_j(y)F_i(y) - f_i(y)F_j(y) \geq 0 \quad \text{for every } y \leq c.$$

The remaining product term is also nonnegative, so $p_j - p_i \geq 0$. Since $i < j$ was arbitrary, the argmax probabilities are nondecreasing in the mean parameter:

$$p_1 \leq p_2 \leq \cdots \leq p_M.$$

For the final claim, let B be a subset consisting of the k largest means. Then

$B = \{M - k + 1, \dots, M\}$. Because $\sum_{i=1}^M p_i = 1$, the average of the k largest probabilities is at least the overall average:

$$\frac{1}{k} \sum_{i \in B} p_i \geq \frac{1}{M} \sum_{i=1}^M p_i = \frac{1}{M}.$$

Hence

$$\mathbb{P}(\arg \max_j Y_j^{(c)} \in B) = \sum_{i \in B} p_i \geq \frac{k}{M}.$$

This proves the lemma. □

Lemma A8.43 (Conditional law of current survivors). *Fix $T < \phi$, and let Ω_T^N denote the event that ME remains active and affordability is inactive through Green turn $\lfloor TN \rfloor$. For Green turn $t \leq \lfloor TN \rfloor$, let S_t be the set of plots available immediately before the Green pick, and set $n_t := |S_t| = N - 2t + 1$. Define*

$$\mathcal{G}_t := \sigma(X_1, \dots, X_N, (g_s, Y_{g_s})_{1 \leq s < t}).$$

Then S_t , n_t , and the surviving latent agricultural profile $(X_i)_{i \in S_t}$ are \mathcal{G}_t -measurable. Moreover, if

$$c_1 := +\infty, \quad c_t := Y_{g_{t-1}} \quad (t \geq 2),$$

then, conditional on \mathcal{G}_t , the current survivor scores $(Y_i)_{i \in S_t}$ are independent and each has law

$$Y_i \mid \mathcal{G}_t \sim \mathcal{N}(\rho X_i, 1 - \rho^2) \quad \text{truncated to } (-\infty, c_t).$$

Proof. Conditional on the full latent agricultural profile $\mathbf{X} = (X_1, \dots, X_N)$, the ecological latent scores are independent with

$$Y_i \mid X_i = x_i \sim \mathcal{N}(\rho x_i, 1 - \rho^2).$$

Farmer deletions depend only on \mathbf{X} and the past Green indices, so S_t and the current surviving X -profile are \mathcal{G}_t -measurable.

On Ω_T^N , Green always removes the current $\arg \max Y$. Since the latent law is continuous and the survivor set shrinks over time,

$$Y_{g_1} > Y_{g_2} > \cdots > Y_{g_{t-1}} \quad \text{a.s.}$$

Hence

$$c_t = Y_{g_{t-1}} = \min_{1 \leq s < t} Y_{g_s} \quad \text{a.s.}$$

If $i \in S_t$, then i survived every earlier Green turn $s < t$, so necessarily

$$Y_i < Y_{g_s} \quad \text{for all } s < t.$$

Because the Green-picked values are strictly decreasing, this is equivalent to the single coordinatewise constraint

$$Y_i < c_t.$$

Thus, conditional on \mathcal{G}_t , each current survivor $i \in S_t$ faces only the one-coordinate upper truncation $Y_i < c_t$. Integrating out the coordinates outside S_t preserves factorization across S_t , so the conditional joint density factors as

$$f((y_i)_{i \in S_t} \mid \mathcal{G}_t) \propto \prod_{i \in S_t} \varphi_{\rho X_i, 1-\rho^2}(y_i) \mathbf{1}\{y_i < c_t\}.$$

That is exactly the product law of independent upper-truncated normals with means ρX_i and common variance $1 - \rho^2$. □

Proposition A8.44 (Conditional top-block hit bound). *Fix $T < \phi$ and work on Ω_T^N . For Green turn $t \leq \lfloor TN \rfloor$, let $B_t(k) \subseteq S_t$ be the indices of the k highest surviving X -values. Then*

$$(21) \quad \mathbb{P}(g_t \in B_t(k) \mid \mathcal{G}_t) \geq \frac{k}{n_t} \quad \text{a.s.}$$

In particular, if B_α^N is the initial top- $\lfloor \alpha N \rfloor$ A -block and $X_t^{N,\alpha}$ denotes the number of its

survivors in S_t , then

$$(22) \quad \mathbb{P}\left(g_t \in B_\alpha^N \mid \mathcal{G}_t\right) \geq \frac{X_t^{N,\alpha}}{n_t} \quad a.s.$$

whenever $X_t^{N,\alpha} > 0$.

Proof. By Lemma A8.43, conditional on \mathcal{G}_t , the current survivor scores are independent upper-truncated normals with common upper truncation c_t , common variance $1 - \rho^2$, and means ρX_i , which are ordered exactly as the surviving X_i 's because $\rho > 0$. Therefore the hypotheses of Lemma A8.42 are satisfied conditional on \mathcal{G}_t . That lemma shows that argmax probabilities are nondecreasing in the mean parameter, hence nondecreasing in X_i . Summing the k largest of these n_t probabilities yields

$$\mathbb{P}(g_t \in B_t(k) \mid \mathcal{G}_t) \geq \frac{k}{n_t}.$$

For the fixed initial block B_α^N , its current survivors are precisely the $X_t^{N,\alpha}$ highest surviving X -values, so (22) is the special case $k = X_t^{N,\alpha}$. \square

Proposition A8.45 (Correlated Farmer-frontier lower envelope). *Let*

$$\alpha_\star(s) := 1 - \sqrt{1 - 2s}, \quad 0 \leq s < 1/2.$$

Fix $T < \phi$. Then for every $\varepsilon > 0$,

$$(23) \quad \mathbb{P}\left(\Omega_T^N \cap \left\{ \inf_{0 \leq s \leq T} \frac{F_{\lfloor sN \rfloor}^{N,ME}}{N} < \alpha_\star(s) - \varepsilon \right\}\right) \rightarrow 0.$$

Consequently, whenever $\mathbb{P}(\Omega_T^N) \rightarrow 1$, the unconditional frontier envelope

$$\mathbb{P}\left(\inf_{0 \leq s \leq T} \frac{F_{\lfloor sN \rfloor}^{N,ME}}{N} \geq \alpha_\star(s) - \varepsilon\right) \rightarrow 1$$

follows.

Proof. Fix $\alpha \in (0, 1)$, and let B_α^N be the initial top- $\lfloor \alpha N \rfloor$ A -block. For Green turn t , write $X_t^{N,\alpha}$ for the number of members of B_α^N still present in S_t . On Ω_T^N , the Green rule is exact

argmax- Y , so by Proposition A8.44,

$$p_t^\alpha := \mathbb{P}\left(g_t \in B_\alpha^N \mid \mathcal{G}_t\right) \geq \frac{X_t^{N,\alpha}}{n_t} \quad \text{a.s.}$$

Couple $X_t^{N,\alpha}$ to a baseline comparison chain by adjoining an i.i.d. sequence $(U_t)_{t \geq 1}$ of $\text{Unif}[0, 1]$ variables, independent of the latent draws. Realize the actual Green hit indicator on B_α^N as

$$G_t^\alpha := \mathbf{1}\{U_t \leq p_t^\alpha\}.$$

Then, on Ω_T^N ,

$$X_{t+1}^{N,\alpha} = (X_t^{N,\alpha} - 1 - G_t^\alpha)^+,$$

because after Green's pick, the next Farmer deletion removes one further block item whenever any block item remains.

Now define the comparison chain $\tilde{X}_t^{N,\alpha}$ by

$$\tilde{X}_1^{N,\alpha} := X_1^{N,\alpha}, \quad \tilde{X}_{t+1}^{N,\alpha} := \left(\tilde{X}_t^{N,\alpha} - 1 - \mathbf{1}\left\{U_t \leq \frac{\tilde{X}_t^{N,\alpha}}{n_t}\right\}\right)^+.$$

This is the baseline fixed-block depletion chain with uniform Green hits (as established by the dynamic survivor law, Lemma A8.49).

Let

$$f_t(k, u) := \left(k - 1 - \mathbf{1}\left\{u \leq \frac{k}{n_t}\right\}\right)^+, \quad k \in \{0, 1, \dots, n_t\}.$$

For each fixed u , the map $k \mapsto f_t(k, u)$ is nondecreasing. Since $p_t^\alpha \geq X_t^{N,\alpha}/n_t$, we have

$$X_{t+1}^{N,\alpha} = (X_t^{N,\alpha} - 1 - \mathbf{1}\{U_t \leq p_t^\alpha\})^+ \leq f_t(X_t^{N,\alpha}, U_t).$$

If $X_t^{N,\alpha} \leq \tilde{X}_t^{N,\alpha}$, monotonicity of f_t gives

$$f_t(X_t^{N,\alpha}, U_t) \leq f_t(\tilde{X}_t^{N,\alpha}, U_t) = \tilde{X}_{t+1}^{N,\alpha}.$$

Since $X_1^{N,\alpha} = \tilde{X}_1^{N,\alpha}$, induction yields the pathwise domination

$$(24) \quad X_t^{N,\alpha} \leq \tilde{X}_t^{N,\alpha} \quad \text{for all } t \leq \lfloor TN \rfloor \quad \text{on } \Omega_T^N.$$

The scaled comparison chain

$$\tilde{x}_N^\alpha(s) := \frac{1}{N} \tilde{X}_{\lfloor sN \rfloor}^{N,\alpha}$$

is exactly the baseline fixed-block depletion chain established by the dynamic survivor law (Lemma A8.49): its jumps are bounded by $2/N$, and its limiting drift on the positive region is

$$f(x, s) = -1 - \frac{x}{1 - 2s}.$$

Hence the hypotheses of Theorem A8.28 are satisfied, so Proposition A8.29 and Corollary A8.30 apply unchanged to $\tilde{X}_t^{N,\alpha}$. Therefore the comparison chain has fluid limit

$$x'_\alpha(s) = -1 - \frac{x_\alpha(s)}{1 - 2s}, \quad x_\alpha(0) = \alpha,$$

with solution

$$x_\alpha(s) = (1 - 2s) + (\alpha - 1)\sqrt{1 - 2s}.$$

The conversion from fixed-block depletion to frontier exhaustion is then the same as in the baseline support appendix. Fix $\varepsilon > 0$, choose a grid

$$0 = u_0 < u_1 < \cdots < u_m = T$$

so fine that

$$\alpha_\star(u_j) - \alpha_\star(u_{j-1}) < \varepsilon/2 \quad \text{for all } j,$$

and define

$$\alpha_j := \max\{0, \alpha_\star(u_j) - \varepsilon/2\}.$$

Since $\alpha_j < \alpha_\star(u_j)$, the deterministic comparison path x_{α_j} hits zero strictly before u_j . The baseline LLN for \tilde{X}^{N,α_j} , together with the pathwise domination (24), implies that on Ω_T^N the actual initial top- $\alpha_j N$ block is extinct by time $u_j N$ with probability tending to one.

Extinction of that block is exactly the event

$$\frac{F_{\lfloor u_j N \rfloor}^{N,\text{ME}}}{N} \geq \alpha_j.$$

A union bound over the finite grid and monotonicity of $t \mapsto F_t^{N,\text{ME}}$ then yield (23). \square

A8.6.4 Retiring the host-depletion turn-count certificate

Earlier drafts tried to prove the $\rho = 0.3$ turn-count lower bound by charging continuation purchases to a static host pool. That route is not used in the proof below. Under positive correlation, static host-depletion bookkeeping has the wrong monotonicity for a continuation-cost certificate: removing high- Y items from a host can make the remaining top- Y tier cheaper or more expensive depending on the deletion pattern. The final proof therefore uses the dynamic survivor law and the fluid budget path directly.

The only deterministic numerical input needed for affordability is the ODE cost certificate recorded in Section A8.6.6: along the validated fluid trajectory, Green's cumulative agricultural spend through $\kappa_{\text{low}} = 0.620816$ is at most

$$2.371402296351 < 2.525.$$

Thus the fluid Green budget has the explicit floor

$$\beta_0 := 2.525 - 2.371402296351 > 0.153597.$$

Proposition A8.60 converts this deterministic budget floor into a finite- N affordability statement by a stopping-time bootstrap. No static continuation-host depletion inequality is used.

Retained economic observation. The following pathwise observation records the economic content of a naïve max- A Farmer. It is retained for interpretation and for companion-paper references, but it is not used in the stopping-time affordability bootstrap or in Proposition A8.60.

Lemma A8.46 (Farmer deletion cannot increase future top- Y cost). *Let H be a finite set of available plots. Order H by descending Y -value:*

$$h_{(1)}^Y, h_{(2)}^Y, \dots, h_{(|H|)}^Y.$$

For $v \leq |H|$, define the cost of the next v Green purchases from H under \max - Y selection:

$$T_v(H) := \sum_{\ell=1}^v A(h_{(\ell)}^Y).$$

If $j^* \in H$ satisfies $A(j^*) = \max_{h \in H} A(h)$, then for every $v \in \{1, \dots, |H| - 1\}$,

$$T_v(H \setminus \{j^*\}) \leq T_v(H).$$

Proof. If j^* is not among the top- v Y -ranked items of H , then deleting it leaves the top- v Y -set unchanged, so T_v is unchanged.

If j^* is among the top- v Y -items, then after deletion the new top- v Y -set is obtained by removing j^* and promoting the item previously at Y -rank $v + 1$. Therefore

$$T_v(H \setminus \{j^*\}) = T_v(H) - A(j^*) + A(h_{(v+1)}^Y).$$

Since j^* has maximal A in all of H ,

$$A(j^*) \geq A(h_{(v+1)}^Y),$$

hence $T_v(H \setminus \{j^*\}) \leq T_v(H)$. □

Remark A8.47 (Economic content). The lemma captures a structural feature of adversarial procurement under a naïve \max - A Farmer. When Farmer removes items from a host pool, it always removes the most agriculturally expensive item. If that item was one that Green intended to purchase (because it also had high ecological value), Green substitutes a cheaper alternative—cheaper because the removed item was the single most expensive item in the entire pool. If Farmer’s removal targets an item Green did not intend to purchase, Green is unaffected. In either case, Farmer’s interference cannot increase Green’s future spending. This is the game-theoretic content that an abstract set-monotonicity argument would lack: it is not true that *arbitrary* deletions lower Green’s cost, but it is true that *max*- A deletions do, because the \max - A item is at least as expensive as any substitute Green might need.

The next monotonicity lemma is likewise retained only as economic intuition for truncated hosts under positive correlation. It is not invoked by the affordability bootstrap.

Lemma A8.48 (Monotonicity of conditional cost on a truncated host). *Fix $q \in (0.1, 10)$, let $x_q := \Phi^{-1}((q - 0.1)/9.9)$, and define*

$$m_q(y) := \mathbb{E}[A \mid Y = y, A \leq q].$$

Under the Gaussian copula with $\rho > 0$, the map $y \mapsto m_q(y)$ is strictly increasing on \mathbb{R} .

Proof. Write $A = 0.1 + 9.9\Phi(X)$, so $\{A \leq q\} = \{X \leq x_q\}$. Then

$$m_q(y) = 0.1 + 9.9 \mathbb{E}[\Phi(X) \mid Y = y, X \leq x_q].$$

Since Φ is increasing, it suffices to show that the conditional distribution of X given $(Y = y, X \leq x_q)$ is stochastically increasing in y .

Under the copula, $X \mid Y = y \sim N(\rho y, \sigma^2)$ with $\sigma^2 = 1 - \rho^2$. The density of $X \mid (Y = y, X \leq x_q)$ is therefore

$$f_{X|Y=y, X \leq x_q}(x) = \frac{\varphi_\sigma(x - \rho y)}{\Phi_\sigma(x_q - \rho y)} \mathbf{1}\{x \leq x_q\},$$

where φ_σ and Φ_σ are the $N(0, \sigma^2)$ density and cdf. This is a location family truncated above at a fixed point. The untruncated family $N(\rho y, \sigma^2)$ has monotone likelihood ratio in y (Gaussian shift family). Truncation to $(-\infty, x_q]$ preserves the MLR property, because the truncation set does not depend on y . Stochastic ordering follows from MLR ordering, so $\mathbb{E}[\Phi(X) \mid Y = y, X \leq x_q]$ is increasing in y . \square

A8.6.5 Exact dynamic survivor law

The conditional survivor law in Lemma A8.43 conditions on the full latent agricultural profile \mathbf{X} and shows that the Y_i 's are independent truncated normals with heterogeneous means ρX_i . For the frontier envelope, this suffices. For the fluid-limit state reduction in both the baseline and correlated theorems, a stronger result is needed: the survivor pool must be iid from the doubly truncated host $\nu_{x,c}$ conditional *only* on the current thresholds (x_t, c_{t-1}) , not on the full X -profile.

Lemma A8.49 (Exact dynamic survivor law). *Fix $\rho \in [0, 1]$ and consider the theorem object on any event on which Farmer and Green both remain affordable, so they play exact $\arg \max X$ and $\arg \max Y$. Let i_t^F and i_t^G denote the indices of Farmer's and Green's t -th purchases; write $x_t := X_{i_t^F}$ and $c_t := Y_{i_t^G}$, with the convention $c_0 := +\infty$. Let*

$$\mathcal{H}_t := \sigma(i_s^F, X_{i_s^F}, Y_{i_s^F} : s \leq t; \quad i_{s'}^G, X_{i_{s'}^G}, Y_{i_{s'}^G} : s' \leq t-1)$$

denote the history visible immediately before Green's t -th move. Because the Gaussian-copula law is continuous, ties and boundary coincidences have probability zero; all statements below are therefore unchanged if strict inequalities are replaced by weak inequalities. Conditional on \mathcal{H}_t , the number of surviving labels is deterministic, namely $N - 2t + 1$ before Green's t -th move, and the regular conditional law of the surviving sample is iid from $\nu_{x_t, c_{t-1}}$: the Gaussian-copula law restricted to $\{X \leq x_t\} \cap \{Y \leq c_{t-1}\}$.

Proof. The argument rests on the following two-dimensional analogue of the standard order-statistics factorisation. It is stated in density language for readability; equivalently it is a statement about regular conditional probabilities, and the continuity assumption removes all tie and boundary events.

Factorisation sublemma. *Let $(Z_1, W_1), \dots, (Z_n, W_n)$ be iid from a joint law P on \mathbb{R}^2 with continuous Z -marginal, and let $k^* := \arg \max_j Z_j$ (a.s. unique). Then, conditional on $\{k^* = k, Z_k = z, W_k = w\}$, the remaining pairs $(Z_j, W_j)_{j \neq k}$ are iid from P restricted to $\{Z < z\}$, independently of w .*

Proof of sublemma. By exchangeability, k^* is uniform on $\{1, \dots, n\}$. Condition on $k^* = k$ and fix the value $(Z_k, W_k) = (z, w)$. The remaining $n - 1$ pairs are iid from P subject to the single event $\{\max_{j \neq k} Z_j < z\}$. The conditional joint density of $(Z_j, W_j)_{j \neq k}$ given this event is

$$\frac{\prod_{j \neq k} p(z_j, w_j) \mathbf{1}\{z_j < z\}}{P(\{Z < z\})^{n-1}} = \prod_{j \neq k} p_{|Z < z}(z_j, w_j),$$

where $p_{|Z < z}$ is the density of P truncated to $\{Z < z\}$. The right-hand side does not involve w , so the remaining pairs are iid from $P_{|Z < z}$ and independent of w . \square

We now prove the lemma by induction on t .

Base case ($t = 1$). The original N pairs are iid from the Gaussian copula. Farmer's first pick is $\arg \max X$; applying the factorisation sublemma with $Z = X$, $W = Y$, $P = \text{copula law}$, the remaining $N - 1$ pairs are iid from the copula law restricted to $\{X < x_1\}$ conditional on \mathcal{H}_1 . Because $c_0 = +\infty$, this is ν_{x_1, c_0} .

Inductive step. Assume the conclusion holds before Green turn t : conditional on \mathcal{H}_t , survivors are iid from $\nu_{x_t, c_{t-1}}$. Green picks $\arg \max Y$. Applying the sublemma with $Z = Y$, $W = X$ and base law $\nu_{x_t, c_{t-1}}$, the remaining pairs conditional on $\mathcal{H}_t \cup \{i_t^G, X_{i_t^G}, Y_{i_t^G}\}$ are iid from $\nu_{x_t, c_{t-1}}$ restricted to $\{Y < c_t\}$, which is ν_{x_t, c_t} .

Farmer then picks $\arg \max X$. A second application of the sublemma, this time with $Z = X$, $W = Y$ and base law ν_{x_t, c_t} , gives that the remaining pairs conditional on \mathcal{H}_{t+1} are iid from ν_{x_{t+1}, c_t} . This closes the induction. \square

Remark A8.50 (The Markov reduction is strong, not marginal). Lemma A8.49 conditions on the *full* history \mathcal{H}_t , not only on the marginal (x_t, c_{t-1}) . The conditional law of the survivors depends on \mathcal{H}_t only through the sufficient statistic (x_t, c_{t-1}) , so the normalised state $(x_t, c_{t-1}, b_F(t), b_G(t))$ is Markov. This sufficiency — not merely the marginal identification at (x_t, c_{t-1}) — is what the Kurtz density-dependent LLN requires: the one-step transition kernel must depend on the past only through the current state.

A8.6.6 Global fluid-limit value path up to the certified affordable horizon

The correlated extension is closed by tracking the actual max- Y path in latent coordinates. This avoids the singular bookkeeping created by static host-depletion certificates. Throughout this subsection

$$M(x, c; \rho) := \Phi_2(x, c; \rho), \quad \sigma := \sqrt{1 - \rho^2},$$

and φ denotes the standard normal density. We also write

$$a(x) := 0.1 + 9.9 \Phi(x), \quad e(c) := 0.1 + 9.9 \Phi(c).$$

For Green's agricultural cost at a max- Y pick from the doubly truncated host $\{X \leq x, Y \leq c\}$,

define

$$(25) \quad g_A(x, c) := \mathbb{E}[A \mid Y = c, X \leq x].$$

Since $X \mid Y = c \sim N(\rho c, \sigma^2)$, this is equivalently

$$(26) \quad g_A(x, c) = 0.1 + 9.9 \frac{\int_{-\infty}^x \Phi(u) \frac{1}{\sigma} \varphi\left(\frac{u - \rho c}{\sigma}\right) du}{\Phi\left(\frac{x - \rho c}{\sigma}\right)}.$$

A fully explicit bivariate-normal form, useful for numerical evaluation, is

$$(27) \quad g_A(x, c) = 0.1 + 9.9 \frac{\Phi_2\left(\frac{x - \rho c}{\sigma}, \frac{\rho c}{\sqrt{2 - \rho^2}}; -\frac{\sigma}{\sqrt{2 - \rho^2}}\right)}{\Phi\left(\frac{x - \rho c}{\sigma}\right)}.$$

Indeed, the numerator in (26) is $\Pr(W \leq X, X \leq x)$ with W an independent standard normal.

Fluid state and ODE. The normalized state is

$$Z(s) := (x(s), c(s), b_F(s), b_G(s)),$$

where $x(s)$ is the current Farmer frontier on the latent X -scale, $c(s)$ is the current Green cutoff on the latent Y -scale, and b_F, b_G are normalized remaining budgets.

During the Farmer-active phase, $s \in (0, \tau)$, the limiting state satisfies

$$(28) \quad \dot{x}(s) = -\frac{1}{\varphi(x(s)) \Phi\left(\frac{c(s) - \rho x(s)}{\sigma}\right)},$$

$$(29) \quad \dot{c}(s) = -\frac{1}{\varphi(c(s)) \Phi\left(\frac{x(s) - \rho c(s)}{\sigma}\right)},$$

$$(30) \quad \dot{b}_F(s) = -a(x(s)), \quad \dot{b}_G(s) = -g_A(x(s), c(s)).$$

The Farmer-exhaustion time is

$$\tau := \inf\{s > 0 : b_F(s) = 0\}.$$

After Farmer exhaustion, x freezes at $x_\tau := x(\tau)$, and on $s \in (\tau, \kappa_{\text{low}})$,

$$(31) \quad \dot{x}(s) = 0, \quad \dot{c}(s) = -\frac{1}{\varphi(c(s)) \Phi\left(\frac{x_\tau - \rho c(s)}{\sigma}\right)}, \quad \dot{b}_G(s) = -g_A(x_\tau, c(s)).$$

At every affordable Green turn, the per-turn ecological value is

$$v(s) = e(c(s)).$$

Lemma A8.51 (Empirical lower-orthant mass along endogenous frontiers). *On any stopped horizon on which exact max-X/max-Y play holds, let n_t be the number of surviving labels immediately before Green's t -th move and let (x_t, c_{t-1}) be the corresponding frontiers. Then, for every fixed $T < 1/2$,*

$$\sup_{1 \leq t \leq \lfloor TN \rfloor} \left| \frac{n_t}{N} - \Phi_2(x_t, c_{t-1}; \rho) \right| \xrightarrow{p} 0.$$

More precisely, the difference is $O_p(N^{-1/2}) + O(N^{-1})$ uniformly over the stopped path.

Proof. By the pathwise part of Lemma A8.49, after Farmer's t -th purchase and before Green's t -th purchase the surviving labels are exactly the initial labels in the lower orthant

$$\{X \leq x_t, Y \leq c_{t-1}\},$$

apart from at most the just-selected boundary label; ties have probability zero under the continuous copula law. Hence

$$\left| \frac{n_t}{N} - \frac{1}{N} \sum_{i=1}^N \mathbf{1}\{X_i \leq x_t, Y_i \leq c_{t-1}\} \right| \leq \frac{1}{N}.$$

The class of bivariate lower orthants is VC, so its empirical process is Glivenko–Cantelli; in particular,

$$\sup_{x,c} \left| \frac{1}{N} \sum_{i=1}^N \mathbf{1}\{X_i \leq x, Y_i \leq c\} - \Phi_2(x, c; \rho) \right| \xrightarrow{p} 0.$$

Evaluating this uniform bound at the random thresholds (x_t, c_{t-1}) gives the claim. \square

Lemma A8.52 (One-step drift in latent coordinates). *On any compact interior set with $x, c \in \mathbb{R}$ and positive remaining budgets, the one-step conditional drift of the normalized state under exact max-X Farmer and exact max-Y Green play is given by (28)–(30), up to a*

uniform $o_p(1)$ error after the usual time scaling $s = t/N$. After Farmer exhaustion the same argument gives (31) with x frozen.

Proof. By Lemma A8.49, conditional on the history immediately before Green's turn at thresholds (x_t, c_{t-1}) , the survivor pool is iid from the bivariate normal law restricted to $\{X \leq x_t, Y \leq c_{t-1}\}$. Its population mass is $M(x_t, c_{t-1}; \rho)$, and its marginal X -density at the upper frontier is

$$f_X(x_t; x_t, c_{t-1}) = \frac{\varphi(x_t) \Phi\left(\frac{c_{t-1} - \rho x_t}{\sigma}\right)}{M(x_t, c_{t-1}; \rho)}.$$

Lemma A8.51 gives $n_t = NM(x_t, c_{t-1}; \rho) + O_p(\sqrt{N})$ uniformly on the stopped compact interior sets under consideration. The standard endpoint spacing formula for the maximum of n_t iid draws from a density positive at the endpoint gives

$$\mathbb{E}[x_t - x_{t+1} \mid \mathcal{H}_t] = \frac{1}{n_t f_X(x_t; x_t, c_{t-1})} + O(N^{-2}) = \frac{1}{N \varphi(x_t) \Phi\left(\frac{c_{t-1} - \rho x_t}{\sigma}\right)} + O(N^{-3/2}).$$

The factor M cancels between the survivor count and the conditional density normalization. Since one macroscopic time step is $1/N$, this yields the displayed \dot{x} drift. The c -drift is identical with X and Y interchanged:

$$\mathbb{E}[c_{t-1} - c_t \mid \mathcal{H}_t] = \frac{1}{N \varphi(c_{t-1}) \Phi\left(\frac{x_t - \rho c_{t-1}}{\sigma}\right)} + O(N^{-3/2}).$$

Farmer's pick concentrates at $X = x_t$, so its agricultural value is $a(x_t) + o(1)$. Green's max- Y pick concentrates at $Y = c_{t-1}$; conditional on that event, its X -coordinate has law $N(\rho c_{t-1}, \sigma^2)$ truncated to $(-\infty, x_t]$, giving expected agricultural cost $g_A(x_t, c_{t-1}) + o(1)$. These are the budget drifts. Once Farmer is exhausted, no further x -updates occur, and the same max- Y calculation with the frozen frontier gives (31). \square

Remark A8.53 (Tail coordinates). Tail variables are still useful for plotting:

$$u(s) := 1 - \Phi(x(s)), \quad w(s) := 1 - \Phi(c(s)).$$

Whenever (28)–(29) hold,

$$\dot{u}(s) = \frac{1}{\Phi\left(\frac{c(s) - \rho x(s)}{\sigma}\right)}, \quad \dot{w}(s) = \frac{1}{\Phi\left(\frac{x(s) - \rho c(s)}{\sigma}\right)}.$$

These formulas are included only as a coordinate translation; the proof uses latent (x, c) coordinates.

Lemma A8.54 (Tail-coordinate entrance uniqueness). *Fix $\rho \in [0, 0.5]$, set $\sigma = \sqrt{1 - \rho^2}$, and write*

$$q(r) := \Phi^{-1}(1 - r), \quad r \in (0, 1),$$

where Φ is the standard normal distribution function. In tail coordinates

$$u(s) := 1 - \Phi(x(s)), \quad w(s) := 1 - \Phi(c(s)),$$

the Farmer-active frontier equations are

$$(32) \quad \dot{u}(s) = F_\rho(u(s), w(s)) := \frac{1}{\Phi\left(\frac{q(w(s)) - \rho q(u(s))}{\sigma}\right)},$$

$$(33) \quad \dot{w}(s) = G_\rho(u(s), w(s)) := \frac{1}{\Phi\left(\frac{q(u(s)) - \rho q(w(s))}{\sigma}\right)}.$$

The system (32)–(33) has a unique entrance solution on $(0, \tau]$ with

$$\lim_{s \downarrow 0} (u(s), w(s)) = (0, 0).$$

It satisfies the mass identity

$$(34) \quad \Phi_2(q(u(s)), q(w(s)); \rho) = 1 - 2s,$$

and, for every

$$0 < \alpha < \alpha_\rho := \frac{1 - \rho}{1 + \rho},$$

has the entrance expansion

$$(35) \quad u(s) = s + O_\alpha(s^{1+\alpha}), \quad w(s) = s + O_\alpha(s^{1+\alpha}) \quad (s \downarrow 0).$$

In particular $u(s)/s \rightarrow 1$ and $w(s)/s \rightarrow 1$.

Proof. Existence is obtained by approximation from the interior. Let $s_n \downarrow 0$, and choose $r_n \in (0, 1/2)$ so that

$$\Phi_2(q(r_n), q(r_n); \rho) = 1 - 2s_n.$$

Such r_n exists for all sufficiently large n , and the Gaussian joint-tail bound used below gives $r_n/s_n \rightarrow 1$. For each n , solve the ordinary initial-value problem for (32)–(33) from $(u_n(s_n), w_n(s_n)) = (r_n, r_n)$. The vector field is smooth away from the axes, so the solution exists until the deterministic Farmer-exhaustion time or until it leaves $(0, 1)^2$. The derivative identity below gives $\Phi_2(q(u_n(s)), q(w_n(s)); \rho) = 1 - 2s$ while the solution exists. Consequently $s \leq u_n(s), w_n(s) \leq 2s$ on every sufficiently small entrance interval. On this ratio band the same Mills-ratio estimate used below bounds the vector field by $1 + C_\alpha s^\alpha$, so the family is equicontinuous near zero and locally equicontinuous on compact subintervals of $(0, \tau]$. A diagonal Arzela–Ascoli argument therefore gives a limit (u, w) with entrance value $(0, 0)$. Passing to the limit in the integral equations on every interval $[\delta, T] \subset (0, \tau]$ shows that this limit is a C^1 solution of (32)–(33).

We now identify every entrance solution and then prove uniqueness. First differentiate $M(q(u), q(w); \rho) = \Phi_2(q(u), q(w); \rho)$ along any C^1 solution of (32)–(33). Since $q'(r) = -1/\varphi(q(r))$,

$$\frac{d}{ds} M(q(u), q(w); \rho) = -\Phi \left(\frac{q(w) - \rho q(u)}{\sigma} \right) \dot{u} - \Phi \left(\frac{q(u) - \rho q(w)}{\sigma} \right) \dot{w} = -2.$$

Together with $M(q(u), q(w); \rho) \rightarrow 1$ at the entrance, this proves (34).

Let

$$J(u, w; \rho) := \Pr(X > q(u), Y > q(w)).$$

Then (34) is equivalent to $u + w - J(u, w; \rho) = 2s$. Because $0 \leq J(u, w; \rho) \leq \min\{u, w\}$, every entrance solution satisfies

$$\max\{u(s), w(s)\} \leq 2s.$$

The right sides in (32)–(33) are at least 1, since the denominators are at most 1. Hence

$u(s) \geq s$ and $w(s) \geq s$. Therefore

$$s \leq u(s), w(s) \leq 2s \quad (s \downarrow 0),$$

so the two tail coordinates remain in a fixed ratio band.

On any such ratio band, the standard Gaussian quantile expansion gives

$$q(w) - \rho q(u) = (1 - \rho + o(1))\sqrt{2 \log(1/s)}$$

uniformly as $s \downarrow 0$, and the same bound holds with u and w interchanged. Mills' ratio therefore implies, for every $\alpha < \alpha_\rho$,

$$1 - \Phi\left(\frac{q(w) - \rho q(u)}{\sigma}\right) + 1 - \Phi\left(\frac{q(u) - \rho q(w)}{\sigma}\right) = O_\alpha(s^\alpha)$$

uniformly in the ratio band. Consequently

$$F_\rho(u(s), w(s)) = 1 + O_\alpha(s^\alpha), \quad G_\rho(u(s), w(s)) = 1 + O_\alpha(s^\alpha),$$

and integration from the entrance gives (35).

It remains to note uniqueness. On the ratio band $s \leq u, w \leq 2s$, the same Mills-ratio calculation gives an integrable local Lipschitz bound for the vector field:

$$\|(F_\rho, G_\rho)(u_1, w_1) - (F_\rho, G_\rho)(u_2, w_2)\| \leq C_\alpha s^{\alpha-1} \|(u_1, w_1) - (u_2, w_2)\|.$$

Since $\alpha > 0$, Gronwall's inequality applies on $(0, s_0]$ after letting the lower endpoint tend to 0. Two entrance solutions have difference $o(s)$ at the lower endpoint by (35), so their difference is identically zero on $(0, s_0]$, and then ordinary ODE uniqueness extends the identity through $(0, \tau]$. □

Lemma A8.55 (Uniform endpoint spacings on the entrance band). *Fix $\rho \in [0, 0.5]$ and $0 < \varepsilon < 1/4$. There are constants $C, c > 0$, depending only on ρ and ε , with the following property. Consider any exact-play active-phase history for which the current latent thresholds*

are $x = q(u)$, $c_0 = q(w)$, the tail coordinates obey the entrance ratio band

$$s \leq u, w \leq 2s, \quad 0 < s \leq \varepsilon,$$

and the realised survivor count and lower-orthant mass satisfy

$$\left| \frac{n}{N} - \Phi_2(q(u), q(w); \rho) \right| = o(1)$$

uniformly over the histories under consideration. Let Δu_N be the increase in the X -tail coordinate caused by the next exact max- X Farmer deletion, and let Δw_N be the increase in the Y -tail coordinate caused by the next exact max- Y Green deletion. Then, uniformly on this class of histories,

$$(36) \quad \mathbb{E}[\Delta u_N \mid \mathcal{H}] = \frac{1}{N \Phi\left(\frac{q(w) - \rho q(u)}{\sigma}\right)} + o(N^{-1}),$$

$$(37) \quad \mathbb{E}[\Delta w_N \mid \mathcal{H}] = \frac{1}{N \Phi\left(\frac{q(u) - \rho q(w)}{\sigma}\right)} + o(N^{-1}).$$

Moreover, for every $r \geq 0$,

$$(38) \quad \Pr(N\Delta u_N > r \mid \mathcal{H}) + \Pr(N\Delta w_N > r \mid \mathcal{H}) \leq Ce^{-cr}.$$

The constants do not depend on N , s , or the realised history.

Proof. We prove the max- X statement; the max- Y statement is symmetric. By Lemma A8.49, conditional on the full history, the surviving labels are distributed as iid draws from the Gaussian-copula law restricted to $\{X \leq q(u), Y \leq q(w)\}$, with survivor count n . In the global X -tail coordinate $r = 1 - \Phi(X)$, the conditional density at the lower endpoint $r = u$ is

$$h_{u,w}(u) = \frac{\Phi((q(w) - \rho q(u))/\sigma)}{\Phi_2(q(u), q(w); \rho)}.$$

The entrance band and $\rho \leq 0.5$ imply that the numerator above is bounded below by a positive constant depending only on (ρ, ε) . Indeed, $q(w) - \rho q(u)$ is uniformly bounded below on $s \leq u, w \leq 2s, 0 < s \leq \varepsilon$, and tends to $+\infty$ as $s \downarrow 0$. The denominator is bounded between

positive constants on the same band for sufficiently small ε .

Let $R_{(1)}$ be the minimum of the n conditional tail coordinates; then $\Delta u_N = R_{(1)} - u$. The conditional probability of the strip $u < R \leq u + z$ is

$$\frac{1}{\Phi_2(q(u), q(w); \rho)} \int_u^{u+z} \Phi\left(\frac{q(w) - \rho q(r)}{\sigma}\right) dr.$$

For $z = O(N^{-1} \log N)$, the integrand equals its endpoint value plus a uniform $o(1)$ relative error on the entrance band; outside this range the exponential bound below makes the contribution to the mean negligible. Hence standard endpoint-order-statistic asymptotics for this triangular array give

$$\mathbb{E}[\Delta u_N | \mathcal{H}] = \frac{\Phi_2(q(u), q(w); \rho)}{n \Phi((q(w) - \rho q(u))/\sigma)} + o(N^{-1}).$$

Using $n/N = \Phi_2(q(u), q(w); \rho) + o(1)$ yields (36).

For the tail bound, the same strip formula and the lower bound on the endpoint integrand imply that, for all $z \geq 0$, the conditional strip mass of $u < R \leq u + z$ is at least $d \min\{z, \varepsilon\}$ for a constant $d = d(\rho, \varepsilon) > 0$, after decreasing d if necessary. Therefore, for $z = r/N$ and $0 \leq r \leq \varepsilon N$,

$$\Pr(\Delta u_N > r/N | \mathcal{H}) \leq (1 - dr/N)^n \leq e^{-c_1 r}.$$

If $r > \varepsilon N$, then the strip of width ε already has conditional mass at least $d\varepsilon$, so the probability is at most $(1 - d\varepsilon)^n \leq e^{-c_2 N}$, which is bounded by $C_1 e^{-c_3 r}$ for $r \leq N$, while the event is impossible for $r > N$. This proves the max- X half of (38); the max- Y half is identical with u and w interchanged. \square

Proposition A8.56 (Small-share launch in tail coordinates). *Fix $\rho \in [0, 0.5]$. Let*

$$\widehat{Z}_N^\rho(s) := (u_N^\rho(s), w_N^\rho(s), b_{F,N}^\rho(s), b_{G,N}^\rho(s))$$

be the finite-game tail-budget state, and let

$$\widehat{Z}^\rho(s) := (u^\rho(s), w^\rho(s), b_F^\rho(s), b_G^\rho(s))$$

be the tail-coordinate fluid solution generated by (32)–(33) and the budget equations in (30).

For every $\eta > 0$,

$$(39) \quad \lim_{\varepsilon \downarrow 0} \limsup_{N \rightarrow \infty} \Pr\left(\|\widehat{Z}_N^\rho(\varepsilon) - \widehat{Z}^\rho(\varepsilon)\| > \eta\right) = 0.$$

For every fixed positive launch share ε , the same conclusion holds for the corresponding latent-budget state $Z_N^\rho = (x_N^\rho, c_N^\rho, b_{F,N}^\rho, b_{G,N}^\rho)$ after applying $x = q(u)$ and $c = q(w)$. This is the launch input used by the latent-coordinate convergence argument below.

Proof. The proof is direct in the fixed- ρ tail coordinates and does not compare the game to the independent case.

Step 1: tail-coordinate drift. The latent-coordinate drift in Lemma A8.52, combined with the chain rule

$$\dot{u} = -\varphi(x)\dot{x}, \quad \dot{w} = -\varphi(c)\dot{c},$$

gives exactly (32)–(33) with $x = q(u)$ and $c = q(w)$. The budget coordinates retain the bounded drifts $\dot{b}_F = -a(q(u))$ and $\dot{b}_G = -g_A(q(u), q(w))$.

Step 2: entrance of the limiting system. Lemma A8.54 gives the unique entrance solution and the rate $u(s), w(s) = s + o(s)$. Since $a(\cdot)$ and $g_A(\cdot, \cdot)$ take values in $[0.1, 10]$, the budgets satisfy

$$b_F^\rho(s) = \mu_A/2 + O(s), \quad b_G^\rho(s) = \mu_A/2 + O(s),$$

so $\widehat{Z}^\rho(s)$ extends continuously to

$$\widehat{Z}^\rho(0) = (0, 0, \mu_A/2, \mu_A/2).$$

Step 3: tightness and subsequential identification. Fix $0 < \varepsilon < \tau$, small enough that $\varepsilon < \mu_A/40$, and view \widehat{Z}_N^ρ as the piecewise-constant interpolation of the tail-budget chain on $[0, \varepsilon]$. Since both initial normalized budgets equal $\mu_A/2 + o_p(1)$ and each player spends at most 10ε over this window, both budgets remain larger than $10/N$ throughout $[0, \varepsilon]$ with probability tending to one. Thus the exact max- X /max- Y survivor law applies throughout the launch interval on this high-probability event.

We first record two path estimates that are uniform on this entrance interval. After $\lfloor sN \rfloor$ Farmer moves, the Farmer frontier cannot exceed the corresponding initial order-statistic frontier, so $u_N^\rho(s) \geq s + o_p(1)$ uniformly on compact subintervals of $(0, \varepsilon]$; the same argument gives $w_N^\rho(s) \geq s + o_p(1)$ for Green. Second, the empirical lower-orthant mass at the realised thresholds is $1 - 2s + o_p(1)$, uniformly along the path, because exactly two labels are removed per active round and lower orthants form a VC class. Since

$M(q(u), q(w); \rho) = 1 - u - w + J(u, w; \rho)$ with $0 \leq J \leq \min\{u, w\}$, this implies the companion upper bound $u_N^\rho(s), w_N^\rho(s) \leq 2s + o_p(1)$. Thus every subsequential limit enters through the same ratio band $s \leq u, w \leq 2s$.

On that band, Lemma A8.55 gives the uniform endpoint-spacing means (36)–(37) and the exponential spacing bound (38). In particular, with probability tending to one, all tail-coordinate jumps on $[0, \varepsilon]$ are at most $N^{-2/3}$, and the martingale part in the Doob decomposition has quadratic variation $O(N^{-1})$. The budget jumps are deterministically bounded by $10/N$. These estimates give tightness of \widehat{Z}_N^ρ in $D([0, \varepsilon])$, and every limit point is continuous and absolutely continuous.

Let $\widehat{Z} = (u, w, b_F, b_G)$ be any subsequential limit. On every interval $[\delta, \varepsilon]$ with $\delta > 0$, the ratio-band bounds put (u, w) in a compact subset of $(0, 1)^2$, where the vector field

$$H_\rho(u, w, b_F, b_G) := (F_\rho(u, w), G_\rho(u, w), -a(q(u)), -g_A(q(u), q(w)))$$

is locally Lipschitz and bounded. The conditional-drift expansion above, the bounded budget drifts, the VC control of the realised lower orthants, and the martingale estimate imply, after passing to the subsequence, that for all $\delta \leq r < s \leq \varepsilon$,

$$\widehat{Z}(s) - \widehat{Z}(r) = \int_r^s H_\rho(\widehat{Z}(t)) dt.$$

Since $\delta > 0$ is arbitrary and the entrance modulus gives $\widehat{Z}(s) \rightarrow (0, 0, \mu_A/2, \mu_A/2)$ as $s \downarrow 0$, every subsequential limit is an entrance solution of the integral equations generated by $(F_\rho, G_\rho, -a, -g_A)$. Lemma A8.54 identifies the only such tail-coordinate entrance solution.

Hence every subsequence has the same limit, and therefore, for each fixed $\varepsilon > 0$,

$$\sup_{s \in [0, \varepsilon]} \|\widehat{Z}_N^\rho(s) - \widehat{Z}^\rho(s)\| \xrightarrow{P} 0.$$

Step 4: the two-scale launch. The fixed- ε convergence from Step 3 implies that the inner lim sup in (39) is zero for every sufficiently small fixed $\varepsilon > 0$. Letting $\varepsilon \downarrow 0$ proves the displayed two-scale launch. If one wants the same positive-share statement at a larger fixed $\varepsilon < \tau$, first launch at a smaller $\varepsilon_0 < \min\{\varepsilon, \mu_A/40\}$ and then apply the ordinary interior density-dependent LLN on $[\varepsilon_0, \varepsilon]$, where the vector field is locally Lipschitz. For the latent state used below, $q(\cdot)$ is smooth on every interval $[u^\rho(\varepsilon)/2, 2u^\rho(\varepsilon)]$ with $\varepsilon > 0$ fixed, so the continuous mapping theorem transfers the positive-share launch from \widehat{Z}_N^ρ to

$$Z_N^\rho = (x_N^\rho, c_N^\rho, b_{F,N}^\rho, b_{G,N}^\rho). \quad \square$$

Proposition A8.57 (Stopped fluid convergence). *Let σ_N^G be the first Green affordability failure after share m :*

$$\sigma_N^G := \inf \left\{ s \geq m : \max_{i \in S_{\lfloor sN \rfloor}} A_i > Nb_G^N(s) \right\},$$

with $\sigma_N^G = +\infty$ if no such failure occurs. Let σ_N^F be the corresponding Farmer affordability time during the active phase,

$$\sigma_N^F := \inf \left\{ s \geq m : \max_{i \in S_{\lfloor sN \rfloor}} A_i > Nb_F^N(s) \right\}.$$

Then, for every fixed $\delta > 0$,

$$\sup_{m \leq s \leq (\kappa_{\text{low}} - \delta) \wedge \sigma_N^G} \|Z_N(s) - Z(s)\| \xrightarrow{P} 0,$$

where the finite state is interpreted after Farmer exhaustion with the realised Farmer frontier frozen. In the active phase the convergence is first proved stopped at $\sigma_N^G \wedge \sigma_N^F$; moreover

$$\Pr(\sigma_N^F \leq \tau - \delta) \rightarrow 0 \quad \text{for every } \delta > 0.$$

Proof. Choose a launch share $\varepsilon < m$. On the event $\{s < \sigma_N^G \wedge \sigma_N^F\}$, Green has played exact max- Y and Farmer has played exact max- X throughout the active interval under consideration. Lemma A8.49 and Lemma A8.52 therefore give a density-dependent chain with

locally Lipschitz drift on every compact tube visited by the fluid solution. The small-jump condition required for the semimartingale fluid-limit theorem is as follows: in normalized-budget coordinates, each increment is deterministically bounded by $\nu/N = 10/N$; in the latent frontier coordinates (x, c) , the endpoint-spacing bound on compact interior tubes gives a maximum single-step jump of order $O_p(\log N/N)$ over $O(N)$ turns (Lemma A8.61), and the corresponding tail moments satisfy the Lindeberg small-jump condition with uniform drift expansion on every compact tube. Launched from the positive-share convergence supplied by Proposition A8.56, the semimartingale decomposition gives

$$\sup_{m \leq s \leq (\tau - \delta) \wedge \sigma_N^G \wedge \sigma_N^F} \|Z_N(s) - Z(s)\| \xrightarrow{p} 0.$$

The Farmer stop can be removed on compact subintervals before τ . The fluid Farmer budget satisfies

$$\beta_F(\delta) := \inf_{0 \leq s \leq \tau - \delta} b_F(s) > 0.$$

On the stopped interval just displayed, b_F^N converges uniformly to b_F . If $\sigma_N^F \leq \tau - \delta$, then at σ_N^F some surviving item has price exceeding $Nb_F^N(\sigma_N^F)$; since all prices are at most 10, this forces $b_F^N(\sigma_N^F) < 10/N$. Uniform convergence would instead give $b_F^N(\sigma_N^F) \geq \beta_F(\delta)/2$ with probability tending to one, a contradiction. Hence $\Pr(\sigma_N^F \leq \tau - \delta) \rightarrow 0$, and the active convergence holds up to $(\tau - \delta) \wedge \sigma_N^G$.

Let

$$\tau_N := \frac{1}{N} \inf\{t : \text{Farmer cannot afford the current max-}X \text{ item}\}.$$

The fluid Farmer budget has a transversal zero at τ :

$$\dot{b}_F(\tau) = -a(x(\tau)) \leq -0.1 < 0.$$

The active convergence just proved and the standard first-passage mapping for uniformly convergent monotone budget paths imply $\tau_N \rightarrow \tau$ in probability and $Z_N(\tau_N) \rightarrow Z(\tau)$.

After τ_N , Farmer may still afford cheaper surviving items. At τ_N the residual Farmer budget is below the current max- X price, hence below $\bar{u} = 10$; since every surviving plot costs at least

$\ell = 0.1$, Farmer can make at most $\lceil 10/\ell \rceil = 100$ further purchases. Consequently

$$0 \leq K_N^F(\text{ME}) - N\tau_N \leq 100,$$

and replacing the actual post- τ_N process by the coupled frozen-Farmer Green-only process with Farmer frontier held at $x_{\tau_N}^N$ changes normalized budgets, threshold coordinates, and cumulative ecological value by $O(1/N)$. The post-Farmer LLN applied to this coupled process therefore gives the same limit for the actual game.

The coupled post- τ_N state is the three-dimensional Green-only process $(c, b_G, x_{\tau_N}^N)$ with the Farmer frontier frozen. The post-Farmer drift (31) is locally Lipschitz in (c, b_G) and continuous in the frozen-frontier parameter. A second semimartingale LLN, stopped only at σ_N^G , gives uniform convergence on every compact subinterval of the post-Farmer phase up to $\kappa_{\text{low}} - \delta$. Concatenating the active and post-Farmer pieces proves the proposition. \square

A8.6.7 Validated numerical certificates for the $\rho = 0.3$ fluid-limit

Three downstream steps of the $\rho = 0.3$ extension rely on numerical inequalities for quantities that are not available in closed form: the Farmer-active Green agricultural cost $c_{\text{up}}^{\text{FA}}$, the post-Farmer continuation cost on $[\tau, \kappa_{\text{low}}]$, and the fluid-limit ecological value integral $\int_m^{\kappa_{\text{low}}} e(c(s)) ds$. Each is obtained as an outward-rounded numerical enclosure produced by the Python scripts `rho03_production_certificate.py` (cost side) and `rho03_fluid_limit_certificate.py` (value side), archived at Zenodo DOI [10.5281/zenodo.19598799](https://doi.org/10.5281/zenodo.19598799). This subsection states the fluid-limit ODE, the enclosure discipline implemented by the scripts, and the two certificate lemmas invoked in the rest of the appendix.

Fluid-limit ODE in tail coordinates. Write

$$q(r) := \Phi^{-1}(1 - r), \quad r \in (0, 1),$$

for the upper-tail quantile of the standard normal, and

$$u(s) := 1 - \Phi(x(s)), \quad v(s) := 1 - \Phi(c(s)),$$

so that $x = q(u)$ and $c = q(v)$. The Farmer's agricultural cost at the frontier is $a(q(u)) = \nu - d u$ with $d := \nu - \ell = 9.9$, and Green's ecological value at the frontier is $e(q(v)) = \nu - d v$. The active-phase tail-coordinate ODE is

$$(40) \quad \dot{u} = F_1(u, v) := \frac{1}{\Phi((q(v) - \rho q(u))/\sigma)}, \quad \dot{v} = F_2(u, v) := \frac{1}{\Phi((q(u) - \rho q(v))/\sigma)},$$

with $\sigma := \sqrt{1 - \rho^2}$, coupled to the budget drifts

$$(41) \quad \dot{b}_F = -(\nu - d u), \quad \dot{b}_G = -G_A(u, v),$$

where Green's conditional agricultural cost is

$$(42) \quad G_A(u, v) := g_A(q(u), q(v)) = \mathbb{E}[A \mid Y = q(v), X \leq q(u)].$$

Because (X, Y) is bivariate normal with $X \mid Y = q(v) \sim N(\rho q(v), \sigma^2)$, (42) admits the explicit representation

$$(43) \quad G_A(u, v) = \ell + d \cdot \frac{\int_{-\infty}^{q(u)} \Phi(x) \cdot \frac{1}{\sigma} \varphi\left(\frac{x - \rho q(v)}{\sigma}\right) dx}{\Phi\left(\frac{q(u) - \rho q(v)}{\sigma}\right)},$$

obtained by marginalizing the bivariate normal joint against the conditional normal density of $X \mid Y = q(v)$ and mapping the standard-normal tail probability $\Phi(x)$ to the uniform marginal $A = \ell + d \cdot \Phi(x)$. Initial conditions are $u(0) = v(0) = 0$ and $b_F(0) = b_G(0) = \mu_A/2 = 2.525$. At the Farmer-exhaustion time $\tau := \inf\{s \geq 0 : b_F(s) \leq 0\}$, u freezes and only (v, b_G) continue to evolve, driven by F_2 and G_A with the frozen- u interval held fixed. The ecological-value integrand used in the value certificate is separately $e(q(v)) = \nu - d v$.

Note that the Green budget drift is $-G_A(u, v)$, not $-e(q(v))$: Green spends agricultural units per purchase (the A-value of the item selected), not ecological units. Because Green plays $\arg \max Y$ over the Farmer-truncated host $\{X \leq x\}$, its per-step cost is the conditional expectation of A given $Y = q(v)$ and $X \leq q(u)$, which is exactly $G_A(u, v)$. The

ecological-value integrand $e(q(v))$ enters the value certificate separately.

Enclosure discipline. Both scripts implement the following enclosure discipline:

- (CI1) **Arithmetic.** IEEE 754 double precision, with outward rounding via `math.nextafter` applied to every arithmetic operation used in an enclosure; a constant outward pad $\varepsilon_{\text{pad}} = 10^{-12}$ absorbs residual IEEE rounding.
- (CI2) **Normal-CDF enclosures.** One-sided bounds for $\Phi(x)$ built from the Abramowitz–Stegun rational approximation (formula 26.2.17), whose published uniform error is $\varepsilon_{\Phi} \leq 7.5 \times 10^{-8}$; inverse-normal brackets $[\Phi_{\text{lo}}^{-1}(p), \Phi_{\text{hi}}^{-1}(p)]$ constructed by verified binary search against those Abramowitz–Stegun bounds, padded outward by ε_{pad} .
- (CI3) **Initialization via diagonal bootstrap.** The ODE symmetry $u(s) \equiv v(s)$ on $[0, \tau]$ with equal zero initial conditions reduces, on the diagonal, to a scalar equation $w' = F_{\text{diag}}(w) := F_1(w, w) = F_2(w, w)$ with $F_{\text{diag}} \geq 1$. Starting at $s_0 = 10^{-5}$ with a trial ceiling $U_{\text{trial}} = 2s_0$, the script rigorously verifies $s_0 \bar{F}_{\text{diag}}(U_{\text{trial}}) \leq U_{\text{trial}}$, yielding the initial box $u(s_0), v(s_0) \in [s_0, s_0 \bar{F}_{\text{diag}}(U_{\text{trial}})]$. This sidesteps the non-Lipschitz corner at the origin where $\Phi^{-1}(1)$ is undefined. The budget and cost accumulators over $[0, s_0]$ are initialised by the pathwise envelopes $-\nu \leq \dot{b}_F, \dot{b}_G \leq 0$ and $0 \leq \dot{c}_G \leq \nu$, giving

$$b_F(s_0), b_G(s_0) \in [\mu_A/2 - \nu s_0, \mu_A/2], \quad c_G(s_0) \in [0, \nu s_0].$$

Since $s_0 < m$, the outside-prefix value accumulator is initialised at $v_{\text{out}}(s_0) = 0$.

- (CI4) **Validated Picard step-tube enclosure.** The validated state at step n is a product interval box $B_n = [\underline{u}, \bar{u}] \times [\underline{v}, \bar{v}] \times [\underline{b}_F, \bar{b}_F] \times [\underline{b}_G, \bar{b}_G] \times [\underline{c}_G, \bar{c}_G] \times [\underline{v}_{\text{out}}, \bar{v}_{\text{out}}]$, extended to include the Green cost accumulator $c_G(s) = \int_0^s G_A(u(r), v(r)) dr$ and the outside-prefix ecological value accumulator

$$v_{\text{out}}(s) := \int_m^{s \vee m} e(q(v(r))) dr,$$

so that $v_{\text{out}}(s) = 0$ for $s \leq m$ and $v_{\text{out}}(s) = \int_m^s e(q(v(r))) dr$ for $s \geq m$. On each step of

size $h = 10^{-4}$, the script constructs a step tube $\mathcal{T}_n \supseteq B_n$ and verifies the Picard inclusion

$$(44) \quad B_n + [0, h] F(\mathcal{T}_n) \subseteq \mathcal{T}_n,$$

where F is the vector field assembled from F_1, F_2 , the budget drifts (41), the Green-cost drift G_A , and the value integrand $e(q(v)) = \nu - dv$. *Inclusion is checked by direct interval containment of each coordinate, not by Lipschitz argument.* The initial tube guess inflates B_n coordinate by coordinate using coordinate-specific drift envelopes:

$F_{\max}^{u,v} = 8 \geq \sup_{[0, \kappa_{\text{low}}]} \max(F_1, F_2)$ for the frontier coordinates, and the pathwise bound

$|\dot{b}_F|, |\dot{b}_G|, |\dot{c}_G|, |\dot{v}_{\text{out}}| \leq \nu = 10$ for the budget and accumulator coordinates. If inclusion fails, the tube is replaced by the hull of the current \mathcal{T}_n and the candidate endpoint box, and the drift is re-evaluated on the enlarged tube. Iteration is bounded by

PICARD_MAX_ITER = 20; across all 6209 steps the maximum number of Picard

iterations actually required was one, so the initial tube always sufficed. Monotonicity of

F_1, F_2 and $e(q(v))$ on \mathcal{T}_n is used to compute those interval images by corner evaluation.

The G_A -interval image is computed by the quadrature enclosure in (CI5). Once

inclusion is verified, the endpoint enclosure

$$B_{n+1} = B_n + h F(\mathcal{T}_n)$$

is formed with outward rounding. Total step count $N_{\text{steps}} = 6209$, partitioned as

$N_{\text{active}} = 3127$ active-phase steps, $N_{\text{uncertain}} = 10$ transition steps (see (CI6)), and

$N_{\text{post}} = 3072$ post-Farmer steps.

(CI5) **Green-cost enclosure.** For every interval box $[\underline{u}, \bar{u}] \times [\underline{v}, \bar{v}]$, the script constructs an outward-rounded interval enclosure

$$G_A([\underline{u}, \bar{u}], [\underline{v}, \bar{v}]) \supseteq \{G_A(u, v) : (u, v) \in [\underline{u}, \bar{u}] \times [\underline{v}, \bar{v}]\}.$$

The numerator in the explicit formula (43), an integral against the conditional normal

density of $X | Y = q(v)$, is enclosed by 96-cell interval quadrature with one-sided

Φ -bounds and outward-rounded evaluations of the normal density; the denominator

$\Phi((q(u) - \rho q(v))/\sigma)$ is enclosed by the corresponding one-sided Φ -bounds. The quotient

is formed by interval arithmetic, with the denominator lower endpoint bounded away

from zero on every certified box used in the integration. Within the Picard step the G_A -interval image is evaluated on the current tube \mathcal{T}_n , not on B_n .

(CI6) **Three-phase active-to-post state machine.** The script maintains a discrete phase flag $\in \{\text{active, uncertain, post}\}$ controlling how the frozen- u constraint is enforced across the Farmer-exhaustion transition:

- *active*: $\underline{b}_F > 0$; both F_1 and F_2 are evaluated on \mathcal{T}_n as above.
- *uncertain*: $\underline{b}_F \leq 0 \leq \overline{b}_F$; some admissible paths may have already exhausted Farmer's budget while others have not, so the script uses a hybrid enclosure in which the lower u - and v -endpoints are left at their pre-step values (corresponding to a path that froze at the start of the step) while the upper endpoints follow the active drift (corresponding to a path still moving under F_1 through the end of the step).
- *post*: $\overline{b}_F \leq 0$; all admissible paths have exhausted. The u -interval is frozen at $[\underline{u}_{\text{current}}, \overline{u}_{\text{current}}]$, and only $(v, b_G, c_G, v_{\text{out}})$ continue to evolve under F_2 and G_A with the frozen u -interval.

The phase transitions active \rightarrow uncertain \rightarrow post are triggered respectively when \underline{b}_F and \overline{b}_F first become non-positive. Over 6209 steps the run logged 10 uncertain steps at the transition.

Lemma A8.58 (Cost certificate for the $\rho = 0.3$ fluid-limit). *Under the integration*

specification (CI1)–(CI6), the validated cost accumulator $c_G(\kappa_{\text{low}}) = \int_0^{\kappa_{\text{low}}} G_A(u(s), v(s)) ds$ is enclosed as

$$c_G(\kappa_{\text{low}}) \in [2.353234170771, 2.371402296351],$$

so the total Green agricultural cost up to κ_{low} is

$$\int_0^{\kappa_{\text{low}}} (-\dot{b}_G(s)) ds = \int_0^{\kappa_{\text{low}}} G_A(u(s), v(s)) ds \leq 2.371402296351.$$

Combined with $b_G(0) = \mu_A/2 = 2.525$,

$$\inf_{0 \leq s \leq \kappa_{\text{low}}} b_G(s) = b_G(\kappa_{\text{low}}) \geq 2.525 - 2.371402296351 \geq 0.153597703649 > 0.$$

The Farmer-exhaustion time is enclosed as $\tau \leq 0.31371$, with the frozen- u interval at $s = \kappa_{\text{low}}$ equal to $[0.426936619958, 0.430911323748]$.

Proof. Direct output of `rho03_production_certificate.py` applied to the ODE (40)–(41) augmented with the cost and value accumulators, under (CI1)–(CI6). The script self-verifies at every one of the 6209 steps: (i) the Picard inclusion (44); (ii) that the interval quadrature for G_A returns intervals with a strictly positive denominator lower endpoint; (iii) the diagonal bootstrap condition $s_0 \bar{F}_{\text{diag}}(U_{\text{trial}}) \leq U_{\text{trial}}$ at $s_0 = 10^{-5}$; (iv) the three-phase state machine (CI6) at the Farmer-exhaustion transition, recording $N_{\text{active}} = 3127$, $N_{\text{uncertain}} = 10$, $N_{\text{post}} = 3072$. The displayed numerical values are extracted verbatim from the machine-readable report `rho03_production_certificate.json`. \square

Lemma A8.59 (Value certificate for the $\rho = 0.3$ fluid-limit). *Under the same integration specification (CI1)–(CI6), the validated value accumulator $v_{\text{out}}(\kappa_{\text{low}}) = \int_m^{\kappa_{\text{low}}} e(c(s)) ds$ is enclosed as*

$$\int_m^{\kappa_{\text{low}}} e(c(s)) ds = \int_m^{\kappa_{\text{low}}} (\nu - dv(s)) ds \in [2.562783061818, 2.585467093762].$$

Combined with the MX threshold $\Theta_{0.3} \leq 2.36413$ established in Proposition A8.64,

$$\int_m^{\kappa_{\text{low}}} e(c(s)) ds - \Theta_{0.3} \geq 2.562783061818 - 2.36413 \geq 0.198653061818 > 0.$$

Proof. Direct output of `rho03_fluid_limit_certificate.py` under (CI1)–(CI6), with the same validated state and Picard step-tube as the cost certificate. The integrand $e(c(s)) = \nu - dv(s) = 10 - 9.9v(s)$ is monotonically decreasing in v , so the lower enclosure of the integral is obtained by integrating the lower endpoint of $\nu - d\bar{v}$ over the step tube; only post-prefix steps with $s \geq m$ contribute via the value-overlap factor in (CI4). The three-phase transition (CI6) is invoked at the Farmer-exhaustion step, delivering the Farmer-exhaustion bound $\tau \leq 0.31371$ matching Lemma A8.58 on the shared grid. The displayed numerical values are extracted verbatim from the machine-readable report produced by the script. \square

Proposition A8.60 (Affordability through κ_{low} by stopping-time bootstrap). *Under the $\rho = 0.3$ theorem object with $\kappa_{\text{low}} = 0.620816$, and with σ_N^G denoting the Green affordability*

stopping time from Proposition A8.57,

$$\Pr(\sigma_N^G > \kappa_{\text{low}}) \rightarrow 1.$$

Consequently Green plays exact max-Y throughout $[m, \kappa_{\text{low}}]$ with probability tending to one.

Proof. By Lemma A8.58,

$$\int_0^{\kappa_{\text{low}}} -\dot{b}_G(s) ds \leq 2.371402296351.$$

Since $B_G^0/N \rightarrow \mu_A/2 = 2.525$, the fluid budget satisfies

$$\beta_0 := \inf_{0 \leq s \leq \kappa_{\text{low}}} b_G(s) = b_G(\kappa_{\text{low}}) \geq 2.525 - 2.371402296351 \geq 0.153597.$$

Fix $\delta > 0$ so small that $10\delta < \beta_0/4$. By Proposition A8.57,

$$\sup_{m \leq s \leq (\kappa_{\text{low}} - \delta) \wedge \sigma_N^G} |b_G^N(s) - b_G(s)| \xrightarrow{p} 0.$$

On the event $\{\sigma_N^G \leq \kappa_{\text{low}} - \delta\}$, the definition of σ_N^G implies that some surviving plot has price exceeding $Nb_G^N(\sigma_N^G)$. Since every plot price is at most 10,

$$b_G^N(\sigma_N^G) < 10/N.$$

Uniform convergence up to the stopped time gives $b_G^N(\sigma_N^G) \rightarrow b_G(\sigma_N^G) \geq \beta_0$ in probability on this event, a contradiction. Hence

$$\Pr(\sigma_N^G \leq \kappa_{\text{low}} - \delta) \rightarrow 0.$$

It remains to rule out a first failure in the terminal strip $(\kappa_{\text{low}} - \delta, \kappa_{\text{low}}]$. With probability tending to one, the preceding paragraph gives $\sigma_N^G > \kappa_{\text{low}} - \delta$, and the stopped convergence at $\kappa_{\text{low}} - \delta$ gives

$$b_G^N(\kappa_{\text{low}} - \delta) > \frac{3\beta_0}{4}.$$

Over the next $\delta N + O(1)$ Green turns, total agricultural spend is at most $10\delta N + O(1)$. Thus

throughout the terminal strip,

$$b_G^N(s) > \frac{3\beta_0}{4} - 10\delta - o(1) > \frac{\beta_0}{2}$$

for all sufficiently large N on a high-probability event. Since $Nb_G^N(s) > 10$ there, Green can afford every surviving plot, so no affordability failure can occur in the terminal strip.

Therefore $\Pr(\sigma_N^G > \kappa_{\text{low}}) \rightarrow 1$. □

Lemma A8.61 (Selected endpoint values track the fluid frontier). *On any stopped compact interval $[\alpha, \beta] \subset (0, \kappa_{\text{low}}]$ on which exact max- Y play holds and*

$$\sup_{\alpha \leq s \leq \beta} \|Z_N(s) - Z(s)\| \xrightarrow{p} 0,$$

the selected Green latent values satisfy

$$\max_{\lfloor \alpha N \rfloor \leq t \leq \lfloor \beta N \rfloor} |Y(g_t^{\text{ME}}) - c(t/N)| \xrightarrow{p} 0.$$

Consequently,

$$\frac{1}{N} \sum_{t=\lfloor \alpha N \rfloor}^{\lfloor \beta N \rfloor} E(g_t^{\text{ME}}) - \frac{1}{N} \sum_{t=\lfloor \alpha N \rfloor}^{\lfloor \beta N \rfloor} e(c(t/N)) \xrightarrow{p} 0.$$

Proof. On $[\alpha, \beta]$ the fluid frontier is bounded away from the singular entrance, so $c(s)$ and $x(s)$ range over a compact subset of \mathbb{R} and the conditional endpoint densities are bounded above and below by positive constants. Lemmas A8.49 and A8.51 therefore give the usual uniform endpoint-spacing bound for the max- Y order statistic: for constants $C, c > 0$,

$$\Pr\left(N |Y(g_t^{\text{ME}}) - c_N(t/N)| > r \mid \mathcal{H}_t\right) \leq Ce^{-cr}$$

uniformly over the stopped histories in the compact interval, where $c_N(t/N)$ is the pre-pick finite- N Green frontier. A union bound over $O(N)$ times with $r = A \log N$, A large enough, gives

$$\max_{\lfloor \alpha N \rfloor \leq t \leq \lfloor \beta N \rfloor} |Y(g_t^{\text{ME}}) - c_N(t/N)| \xrightarrow{p} 0.$$

Uniform state convergence gives $\sup_{\alpha \leq s \leq \beta} |c_N(s) - c(s)| \rightarrow 0$ in probability, proving the first display. Since $e(c) = 0.1 + 9.9\Phi(c)$ is uniformly continuous on the compact range of c , the

Cesàro-value statement follows. □

Proposition A8.62 (Fluid-limit convergence up to the certified horizon). *On the $\rho = 0.3$ ME path, for every fixed $\delta > 0$,*

$$\sup_{m \leq s \leq \kappa_{\text{low}} - \delta} \|Z_N(s) - Z(s)\| \xrightarrow{p} 0,$$

and

$$\frac{1}{N} \sum_{t=\lfloor mN \rfloor + 1}^{\lfloor \kappa_{\text{low}} N \rfloor} E(g_t^{\text{ME}}) \xrightarrow{p} \int_m^{\kappa_{\text{low}}} e(c(s)) ds.$$

Proof. By Proposition A8.60, the stopping-time restriction in Proposition A8.57 is vacuous with probability tending to one on $[m, \kappa_{\text{low}}]$, giving uniform state convergence on every compact subinterval $[m, \kappa_{\text{low}} - \delta]$, for each fixed $\delta > 0$. This is the first displayed limit.

For the value-integral limit we do not need closed-endpoint state convergence. Fix $\varepsilon > 0$ and $\delta > 0$. Applying Lemma A8.61 with $\alpha = m$ and $\beta = \kappa_{\text{low}} - \delta$ identifies the selected Green endpoint values with $e(c(t/N))$ in Cesàro mean on the compact subinterval, and the deterministic Riemann-sum convergence gives

$$\frac{1}{N} \sum_{t=\lfloor mN \rfloor + 1}^{\lfloor (\kappa_{\text{low}} - \delta)N \rfloor} E(g_t^{\text{ME}}) \xrightarrow{p} \int_m^{\kappa_{\text{low}} - \delta} e(c(s)) ds.$$

The terminal strip is bounded deterministically on both sides by the per-item ecological envelope $E \in [0.1, \nu] = [0.1, 10]$:

$$0 \leq \frac{1}{N} \sum_{t=\lfloor (\kappa_{\text{low}} - \delta)N \rfloor + 1}^{\lfloor \kappa_{\text{low}} N \rfloor} E(g_t^{\text{ME}}) \leq 10\delta + O(1/N), \quad 0 \leq \int_{\kappa_{\text{low}} - \delta}^{\kappa_{\text{low}}} e(c(s)) ds \leq 10\delta,$$

because each term is bounded pathwise by ν and there are at most

$$\lfloor \kappa_{\text{low}} N \rfloor - \lfloor (\kappa_{\text{low}} - \delta)N \rfloor = \delta N + O(1) \text{ such terms, and } e(c(s)) = \nu - dv(s) \in [\ell, \nu] \subseteq [0, 10].$$

Choosing δ so that $10\delta \leq \varepsilon/3$ and combining with the compact-subinterval convergence yields

$$\left| \frac{1}{N} \sum_{t=\lfloor mN \rfloor + 1}^{\lfloor \kappa_{\text{low}} N \rfloor} E(g_t^{\text{ME}}) - \int_m^{\kappa_{\text{low}}} e(c(s)) ds \right| \leq \varepsilon$$

with probability tending to one. Taking $\varepsilon \downarrow 0$ gives the displayed value-integral limit.

Boundedness of the per-item ecological values provides uniform integrability when

expectations are taken later. \square

Lemma A8.63 (Static partition of the MX bottom Farmer block). *Let $B_\phi := \{R \geq r_\phi(0.3)\}$, where $\Pr(R \geq r_\phi(0.3)) = \phi$, and let*

$$D_{\phi,0.3} := \{A \leq q_F, R < r_\phi(0.3)\}, \quad q_F := Q_A(1 - \phi).$$

Under the safe-frontier event of Proposition A8.36, the finite MX path has the following deterministic set structure up to $o_p(N)$ labels: Green's Farmer-active purchases are the empirical top-ratio block B_ϕ , Farmer's active purchases are the empirical top-A block $\{A > q_F\}$, and the post-Farmer residual Green pool is $D_{\phi,0.3}$. Moreover the population regions satisfy the exact disjoint identity

$$B_\phi \dot{\cup} D_{\phi,0.3} = \{A \leq q_F\}.$$

Proof. The safe-frontier inequality gives $10/r_\phi(0.3) < q_F$. Therefore every point with $R \geq r_\phi(0.3)$ has $A = E/R \leq 10/r_\phi(0.3) < q_F$, so $B_\phi \subset \{A \leq q_F\}$. This proves the population set identity after splitting $\{A \leq q_F\}$ according to whether $R \geq r_\phi(0.3)$ or $R < r_\phi(0.3)$.

For the finite game, Proposition A8.36 shows that Farmer does not touch the initial top-ratio block before the benchmark exhaustion share and that MX purchases that block in static ratio order. The Farmer path therefore coincides with the descending- A path through exhaustion, and the empirical quantile law replaces its random frontier by q_F up to $o_p(N)$ labels. Once Farmer stops, the remaining labels below the Farmer frontier that are not in the already purchased ratio block are exactly the empirical version of $D_{\phi,0.3}$, again up to $o_p(N)$ boundary labels. The bounded leftover count after Green stops is $O(1)$, hence also $o_p(N)$. \square

Proposition A8.64 (MX threshold identity at $\rho = 0.3$). *Under the $\rho = 0.3$ theorem object,*

$$\frac{1}{N} \text{PC}_N^{\text{MX}} \xrightarrow{p} \mathbb{E}[E \cdot \mathbf{1}\{A \leq q_F\}], \quad q_F := Q_A(1 - \phi),$$

where ϕ is the MX Farmer-exhaustion share. Under the baseline marginals with $\rho = 0.3$, the certified integration interval is

$$\mathbb{E}[E \mathbf{1}\{A \leq q_F\}] \in [3.26511, 3.26513].$$

Proof. By Lemma A8.63, the MX Green purchases are, up to $o_p(N)$ labels, the disjoint union of the static top-ratio block B_ϕ and the post-Farmer residual region $D_{\phi,0.3}$. These two regions are exactly $\{A \leq q_F\}$ at the population level. The thresholded LLN for the bounded weight E , applied to the ratio cutoff and the A -quantile cutoff, therefore gives

$$\frac{1}{N} \text{PC}_N^{\text{MX}} \xrightarrow{p} \mathbb{E}[E \mathbf{1}\{A \leq q_F\}].$$

The displayed interval is the certified bivariate-normal integral enclosure for the same expectation under the $\rho = 0.3$ Gaussian-copula law; the threshold boundary has probability zero, so the thresholded LLN applies without boundary correction. \square

Corollary A8.65 (Direct $\rho = 0.3$ theorem reduction via the fluid-limit value path). *Define the outside-prefix MX benchmark*

$$\Theta_{0.3} := \mathbb{E}[E \mathbf{1}\{A \leq q_F\}] - \underline{M}_E(0.10) \leq 3.26513 - 0.901 = 2.36413.$$

If

$$\int_m^{\kappa_{\text{low}}} e(c(s)) ds > \Theta_{0.3},$$

then

$$\liminf_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}[\text{PC}_N^{\text{ME}} - \text{PC}_N^{\text{MX}}] > 0.$$

Proof. The prefix lower envelope gives

$$\frac{1}{N} \sum_{t=1}^{m_N} E(g_t^{\text{ME}}) \geq \underline{M}_E(0.10) + o_p(1).$$

Proposition A8.62 gives the realised outside-prefix value through the certified affordable horizon:

$$\frac{1}{N} \sum_{t=m_N+1}^{\lfloor \kappa_{\text{low}} N \rfloor} E(g_t^{\text{ME}}) \xrightarrow{p} \int_m^{\kappa_{\text{low}}} e(c(s)) ds.$$

Since any additional ME purchases after κ_{low} have nonnegative ecological value,

$$\frac{1}{N} \text{PC}_N^{\text{ME}} \geq \underline{M}_E(0.10) + \frac{1}{N} \sum_{t=m_N+1}^{\lfloor \kappa_{\text{low}} N \rfloor} E(g_t^{\text{ME}}) + o_p(1).$$

On the MX side, Proposition A8.64 gives the total threshold limit

$$\frac{1}{N} \text{PC}_N^{\text{MX}} \xrightarrow{p} \mathbb{E}[E \mathbf{1}\{A \leq q_F\}].$$

Combining the preceding displays gives, in probability,

$$\frac{1}{N} (\text{PC}_N^{\text{ME}} - \text{PC}_N^{\text{MX}}) \geq \int_m^{\kappa_{\text{low}}} e(c(s)) ds - \left(\mathbb{E}[E \mathbf{1}\{A \leq q_F\}] - \underline{M}_E(0.10) \right) + o_p(1).$$

The parenthesised term is $\Theta_{0.3}$. Because per-item E -values lie in $[0.1, 10]$, the normalised sums are uniformly bounded, hence uniformly integrable. Taking expectations and then the lower limit gives

$$\liminf_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}[\text{PC}_N^{\text{ME}} - \text{PC}_N^{\text{MX}}] \geq \int_m^{\kappa_{\text{low}}} e(c(s)) ds - \Theta_{0.3}.$$

If the displayed integral exceeds $\Theta_{0.3}$, the asymptotic gap is positive. \square

Corollary A8.66 (Validated closure of the $\rho = 0.3$ extension). *Fix $m = 0.10$ and*

$\kappa_{\text{low}} = 0.620816$. *The stopping-time affordability bootstrap gives $K_N^{\text{ME}}/N \geq 0.620816 - o_p(1)$.*

The global fluid-limit value path gives

$$\frac{1}{N} \sum_{t=\lfloor mN \rfloor + 1}^{\lfloor \kappa_{\text{low}} N \rfloor} E(g_t^{\text{ME}}) \xrightarrow{p} \int_m^{\kappa_{\text{low}}} e(c(s)) ds.$$

A validated ODE certificate (Lemma A8.59) establishes

$$\int_m^{\kappa_{\text{low}}} e(c(s)) ds \geq 2.562783061818,$$

while $\Theta_{0.3} = \mathbb{E}[E \mathbf{1}\{A \leq q_F\}] - 0.901 \leq 2.36413$, which follows from the interval

$\mathbb{E}[E \mathbf{1}\{A \leq q_F\}] \in [3.26511, 3.26513]$ *in Proposition A8.64. Consequently,*

$$\liminf_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}[\text{PC}_N^{\text{ME}} - \text{PC}_N^{\text{MX}}] \geq 0.198653061818 > 0.$$

Proof. By Proposition A8.60, Green remains affordable through κ_{low} , so Green plays exact $\arg \max Y$ on $[m, \kappa_{\text{low}}]$. The prefix contribution is the explicit lower-envelope term

$$\frac{1}{N} \sum_{t=1}^{mN} E(g_t^{\text{ME}}) \geq \underline{M}_E(0.10) + o_p(1),$$

and Corollary A8.65 subtracts this term from the total MX threshold through the definition of $\Theta_{0.3}$. Proposition A8.62 identifies Green’s normalised ecological value with the fluid-limit integral. The MX threshold is from Proposition A8.64, which establishes

$\mathbb{E}[E \mathbf{1}\{A \leq q_F\}] \leq 3.26513$ and hence $\Theta_{0.3} \leq 2.36413$. Lemma A8.59 confirms the displayed inequality on the fluid-limit integral, with margin ≥ 0.198653061818 over the threshold $\Theta_{0.3} \leq 2.36413$. Because per-item E -values lie in $[0.1, 10]$, the normalised sums $N^{-1}\text{PC}_N^{\text{ME}}$ and $N^{-1}\text{PC}_N^{\text{MX}}$ are uniformly bounded, hence uniformly integrable. The convergence in probability therefore upgrades to convergence of expectations, and the lim inf inequality follows. \square

Remark A8.67 (Rounded corroborative numerics). The rounded display constants

$$G_{0.3}^{\text{pref}} = 0.30878228, \quad V_{X,0.3}^{\text{out}} = 2.67290033,$$

remain useful corroborative numerics from an earlier version of the proof. The fluid-limit architecture does not use them.

Remark A8.68 (Certificate archival). The Picard step-tube certificate scripts (`rho03_production_certificate.py` for the cost-side report and `rho03_fluid_limit_certificate.py` for the value-side report), together with their human-readable text reports and machine-readable JSON summaries, are archived on Zenodo (DOI: 10.5281/zenodo.19598799). Both scripts share the validated state from (CI4) and produce unified cost and value accumulator outputs; the JSON reports from the two scripts contain the same accumulator intervals $c_G \in [2.353234170771, 2.371402296351]$ and $v_{\text{out}} \in [2.562783061818, 2.585467093762]$ on the shared step grid. A replication check for this paper should confirm that the archived JSONs contain these values; earlier band-reduction reports predating the Picard upgrade are not part of the current certificate chain.

A8.6.8 Technical inputs cited without reproof

This subsection is self-contained at the theorem-architecture level, but several standard or previously isolated inputs are cited without full reproof.

- (S1) **Empirical quantiles and thresholded LLNs.** We use almost-sure convergence of empirical quantiles and thresholded sample means for iid draws under continuous laws,

including conditional laws such as $E \mid A \leq q$. These are standard empirical-process facts and are recorded in Appendix A8.5.

- (S2) **Benchmark-split replacements.** Whenever a stopping time τ_N satisfies $\tau_N/N \rightarrow \tau$ in probability and per-turn values are uniformly bounded, replacing τ_N by $\lfloor \tau N \rfloor$ changes the normalized cumulative sum by $o(1)$. We use this repeatedly to move between random and deterministic split points.
- (S3) **Density-dependent fluid limits.** Proposition A8.45 invokes the standard Kurtz/Wormald law of large numbers for density-dependent Markov chains. The specific chain conditions—bounded $O(1/N)$ increments, compact state space on every fixed $[0, T]$, and locally Lipschitz drift away from absorption—are the same as in Appendix A8.5.
- (S4) **Worst-case deletion principle.** The baseline value bounds and the MX threshold identity use Lemma A8.69. For the $\rho = 0.3$ extension, Lemma A8.46 is retained only as economic interpretation of a naïve max- A Farmer; the cost side is closed by the stopping-time bootstrap in Proposition A8.60. The $\rho = 0.3$ value side uses the global fluid-limit path (Proposition A8.62) rather than deletion bounds.

Lemma A8.69 (Worst-case deletion principle). *Let $z_1 \geq z_2 \geq \dots \geq z_M$ be a descending sequence, and let $D \subset \{1, \dots, M\}$ with $|D| = m$. Let $z_1^{(D)} \geq z_2^{(D)} \geq \dots$ be the descending order of $\{z_j : j \notin D\}$. Then for every $r \leq M - m$,*

$$\sum_{j=1}^r z_j^{(D)} \geq \sum_{j=m+1}^{m+r} z_j.$$

Equivalently: among all deletion sets of size m , the future sum of the next r descending order statistics is minimized by deleting the top m items $\{1, \dots, m\}$.

Proof. After deleting D , at most m of the original top $m + r$ items are missing, so at least r survive. Hence the top r survivors are at least as large as items ranked $m + 1$ through $m + r$ in the original ordering. \square

The project-specific work for $\rho = 0.3$ is exactly what Sections A8.6.3 and A8.6.4 contribute: establishing a usable frontier/value path and closing the affordability argument by a validated

ODE budget bound and stopping-time bootstrap.

A8.7 Ecological conditions for the knapsack reversal

The asymptotic theorems require sufficient ecological heterogeneity. To characterise this requirement, define the *ecological selection premium*

$$\text{ESP} := \frac{\bar{E}_{\text{top-}e} - \bar{E}_{\text{top-}r}}{\bar{E}},$$

where $\bar{E}_{\text{top-}e}$ is the mean ecological value of items in the top decile by E , $\bar{E}_{\text{top-}r}$ is the mean ecological value of items in the top decile by E/A , and \bar{E} is the population mean of E .

Compressed- E simulations, in which only the ecological spread is varied while holding the agricultural distribution fixed, show that the exact normalised PC gap changes sign at $\text{ESP} \approx 12.5\%$, while the current proof architecture requires $\text{ESP} \approx 15\%$. At the baseline ($E \sim U[0.1, 10]$), $\text{ESP} \approx 53\%$, well above either threshold.

The mechanism behind the failure under low ESP is value compression, not frontier collapse: stopping shares, safe-frontier margins, and cost-side geometry are essentially unchanged as E is compressed. What disappears is the early quality gap between the two selection rules. When ecological values are nearly homogeneous, choosing by E versus E/A yields similar ecological portfolios, and the self-hedging advantage has nothing to exploit.

A8.7.1 Why the current proof breaks near compressed E

Remark A8.70 (Diagnostic: terminal-tail movement in compressed- E runs). The following calculation is a numerical diagnostic, not a theorem. Here G , L^c , and L^e denote the simulation script's recorded prefix, common-late, and terminal-tail diagnostic components, respectively. Between the baseline and the exact [4, 6] compressed- E run, the recorded diagnostics are

$$\begin{aligned} G_{\text{base}} - G_{[4,6]} &\approx 0.183971, \\ (L_{[4,6]}^c + L_{[4,6]}^e) - (L_{\text{base}}^c + L_{\text{base}}^e) &\approx 0.022408. \end{aligned}$$

Thus about 89% of the deterioration in the exact normalized gap is attributable to the loss of prefix advantage, while about 11% is attributable to larger late losses. Within the late-loss diagnostic,

$$(45) \quad (L_{[4,6]}^c - L_{\text{base}}^c) + (L_{[4,6]}^e - L_{\text{base}}^e) \approx -0.012812 + 0.035219,$$

so the adverse late movement in this diagnostic table comes from the terminal tail L^e , while the common-late component L^c moves in the opposite direction. These constants are arithmetic summaries of the simulation diagnostics; no population convergence claim is attached to this remark.

Remark A8.71 (Why the terminal tail grows). This is economically natural. The turn-count gap changes very little under pure E -compression:

$$K_N^{\text{ME}}/N \approx 0.695, \quad K_N^{\text{MX}}/N \approx 0.704$$

throughout the whole grid. So the number of overtime MX turns hardly moves. What changes is their ecological value. When E is compressed near 5, each extra MX purchase in the tail carries materially more ecology than in the baseline, where late tail purchases live in a much broader value distribution and include many low- E plots.

A8.8 Finite- N discussion and budget concentration

Throughout this subsection we maintain the baseline assumptions used in the benchmark Monte Carlo design: (e_i, a_i) i.i.d. with marginals $U[0.1, 10]$, independence ($\rho = 0$), budget parity ($s = 0.5$), and a naïve Farmer. The point here is interpretive. We prove a valid budget-concentration lemma, then use it to explain why the very small counterexamples above need not be representative of the paper's 10×10 baseline.

Lemma A8.72 (Decay of κ). *Under the baseline assumptions:*

$$(a) \quad N \kappa(N) \xrightarrow{p} \frac{20}{5.05} \approx 3.9604 \quad \text{as } N \rightarrow \infty, \text{ where } \kappa(N) = \max_i a_i / B_g.$$

(b) For any $c_0 > 20/5.05$,

$$P\left(\kappa > \frac{c_0}{N}\right) \leq \exp\left(-\frac{2N(2.525 - 10/c_0)^2}{24.5025}\right).$$

Proof. Part (a): $\max_i a_i \rightarrow 10$ almost surely (the maximum of N i.i.d. uniforms on $[0.1, 10]$). The denominator $B_g = \frac{1}{2} \sum a_i$ satisfies $B_g/N \xrightarrow{a.s.} 5.05/2$ by the strong law of large numbers. Hence $N\kappa(N) = N \max_i a_i / B_g \xrightarrow{P} 10/(5.05/2) = 20/5.05 \approx 3.9604$. Part (b): since $\max_i a_i \leq 10$ deterministically, $\kappa > c_0/N$ requires $B_g < 10N/c_0$. Apply Hoeffding's inequality to $B_g = \frac{1}{2} \sum a_i$, where each $a_i/2 \in [0.05, 5]$. The Hoeffding range proxy is therefore

$$\sum_{i=1}^N (5 - 0.05)^2 = 24.5025 N.$$

Now $E[B_g/N] = 5.05/2 = 2.525$ and the event $B_g < 10N/c_0$ requires $B_g/N < 10/c_0$. For $c_0 > 20/5.05$ we have $10/c_0 < 2.525$, so the deviation is $2.525 - 10/c_0 > 0$ and Hoeffding gives

$$P(B_g/N < 10/c_0) \leq \exp\left(-\frac{2N(2.525 - 10/c_0)^2}{24.5025}\right),$$

which establishes part (b). □

Table A8.2 reports $E[\kappa]$ for the grid sizes used in this paper.

Table A8.2: Expected budget concentration ratio

Grid	N	$E[\kappa]$ (approx.)
2×2	4	0.795
3×3	9	0.397
4×4	16	0.233
10×10	100	0.039
Bolivia board	390	0.010

Remark A8.73 (What Lemma A8.72 does and does not imply). Lemma A8.72 does *not* by itself prove the baseline asymptotic dominance theorem, nor does it close the $\rho = 0.3$ extension. It shows only that, under the benchmark sampling scheme, any single expensive plot absorbs a vanishing share of the Green budget. This weakens the tight-budget mechanism behind the finite counterexamples and helps explain why the 2×2 knapsack reversals are not representative of the paper's larger baseline boards. The full baseline theorem additionally requires the prefix, common-late, and terminal controls summarized in Appendix A8.3 and

proved in Appendices A8.4–A8.5; the $\rho = 0.3$ extension further requires the validated deterministic certificate supplied in Appendix A8.6.

Remark A8.74 (Simulation-supported comparative statics). In the simulations across the range of dependence parameters θ we examined, the PC gap

$$\Delta(\theta) = \mathbb{E}[\text{PC}(\text{MAXENV}) - \text{PC}(\text{MAXEFF})]$$

is positive and increases with θ from independence through moderate positive dependence, consistent with the self-hedging interpretation: stronger positive association makes high- e plots more likely to be expensive and hence more likely to be targeted early by Farmers. At stronger positive dependence the gap is non-monotone: in the Gaussian-copula family the normalised gap is hump-shaped in ρ , peaking near $\rho \approx 0.7$ before declining as cross-damage amplification begins to dominate selection amplification (see the discussion in Section 4.1.1 and Figure 11). We present these as simulation regularities, not as proved comparative-static theorems.

Remark A8.75 (Reading the tiny-grid anomaly). The tiny-grid exercises below show exactly why theorem inflation is dangerous here. On 2×2 boards, tight budgets can make MAX EFFICIENCY outperform MAX ENVIRONMENT. On the reported 3×3 exercises, the ranking flips and aligns with the larger Monte Carlo simulations. The small-grid evidence is therefore useful as a diagnostic for where the ratio-greedy logic can survive, but it is not an asymptotic proof.

A8.9 Tiny-grid backward-induction exercises

Each reported tiny-grid equilibrium is one *subgame-perfect Nash equilibrium* (SPNE) computed by backward induction on the finite extensive-form game under a deterministic tie-breaking convention. Specifically, whenever the mover has more than one action with the same continuation payoff, the code selects the lowest-index available plot. Because (e_i, a_i) are continuously distributed, exact continuation-payoff ties occur with probability zero almost surely; the deterministic selector is imposed only for reproducibility. We do not claim equilibrium uniqueness, and the reported values should therefore be read as reproducible SPNE benchmarks rather than as bounds over all SPNE profiles.

We compute backward-induction solutions on the finite extensive-form game 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. These calculations are informative but limited to very small grids. They should be read as numerical benchmarks rather than as a general equilibrium theorem for larger boards. In *Claims World*, Hot Spot often lies closest to the equilibrium Green payoff on the reported exercises, while Block Farmers performs poorly. In *Budget World*, Max Efficiency can outperform Max Environment on 2×2 grids, but the ranking reverses on the reported 3×3 exercises and lines up with the larger-grid Monte Carlo pattern. Across both Worlds, first-move effects diminish rapidly with dimension in these computations. The tiny-grid results therefore support a cautious interpretation: tight-budget pathologies are real on minimal boards, but the baseline simulation ranking is already visible by 3×3 under the paper's parameterization.

Three patterns are visible. First, HOT SPOT is the closest of the reported heuristics to the equilibrium Green payoff on many of the Claims World exercises. Second, BLOCK FARMERS is the worst-performing heuristic by a wide and growing margin: on the 4×4 grid at $\rho = 0$, the equilibrium outperforms it by over 13 points on average. Third, MAX ENVIRONMENT improves steadily with dimensionality and is second only to HOT SPOT on the larger grids, consistent with the main Monte Carlo pattern. These results make the small-grid computations informative, but they still do not amount to a general theorem for larger boards.

The rightmost column reveals the key Budget World transition. On the 2×2 grid, MAX EFFICIENCY beats MAX ENVIRONMENT in the majority of replications across all ρ , so the tight-budget counterexample is not merely theoretical. On the 3×3 grid this reverses: MAX ENVIRONMENT wins more than half of head-to-head comparisons, with a positive and growing mean PC gap as ρ increases. The simplest reading is that as N grows from 4 to 9, budget concentration falls sharply and the extreme single-item budget shocks that favour ratio-greedy become less important. Across both grid sizes, BLOCK FARMERS and HOT SPOT are dominated by substantially wider margins than either MAX ENVIRONMENT or MAX EFFICIENCY, so the central Budget World distinction is between value-first and ratio-first purchasing.

Table A8.3: 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 \pm 0.80)	1103/2000 (1.76 \pm 2.40)	1953/2000 (5.53 \pm 2.81)	668/2000 (0.58 \pm 1.14)	912/2000 (45.6%)
	-0.4	406/2000 (0.36 \pm 0.96)	850/2000 (0.94 \pm 2.11)	1760/2000 (4.31 \pm 3.52)	1024/2000 (1.32 \pm 1.86)	1052/2000 (52.6%)
	0.0	342/2000 (0.34 \pm 1.02)	775/2000 (0.60 \pm 1.77)	1606/2000 (3.29 \pm 3.42)	1165/2000 (1.82 \pm 2.24)	1141/2000 (57.0%)
	+0.4	265/2000 (0.25 \pm 0.84)	635/2000 (0.28 \pm 1.35)	1364/2000 (2.18 \pm 2.97)	1401/2000 (2.42 \pm 2.47)	1317/2000 (65.8%)
	+0.8	117/2000 (0.08 \pm 0.45)	441/2000 (0.12 \pm 0.77)	943/2000 (0.85 \pm 1.64)	1598/2000 (2.79 \pm 2.49)	1599/2000 (80.0%)
3×3	-0.8	1331/2000 (1.62 \pm 1.87)	1138/2000 (1.25 \pm 2.10)	1990/2000 (9.47 \pm 3.75)	1386/2000 (1.80 \pm 2.01)	278/2000 (13.9%)
	-0.4	1524/2000 (2.88 \pm 2.61)	1005/2000 (0.71 \pm 2.06)	1944/2000 (8.16 \pm 4.61)	1732/2000 (4.38 \pm 3.21)	415/2000 (20.8%)
	0.0	1501/2000 (2.89 \pm 2.65)	958/2000 (0.44 \pm 2.02)	1865/2000 (6.73 \pm 4.83)	1897/2000 (6.07 \pm 3.38)	607/2000 (30.3%)
	+0.4	1380/2000 (2.25 \pm 2.44)	864/2000 (0.20 \pm 1.87)	1748/2000 (4.89 \pm 4.42)	1957/2000 (7.29 \pm 3.36)	950/2000 (47.5%)
	+0.8	931/2000 (0.92 \pm 1.52)	825/2000 (0.07 \pm 1.20)	1620/2000 (2.37 \pm 2.89)	1991/2000 (7.66 \pm 3.06)	1333/2000 (66.6%)
4×4	-0.8	1757/2000 (3.17 \pm 2.45)	1613/2000 (3.16 \pm 3.30)	2000/2000 (20.01 \pm 4.84)	1770/2000 (3.34 \pm 2.59)	118/2000 (5.9%)
	-0.4	1905/2000 (5.79 \pm 3.22)	1243/2000 (1.15 \pm 3.08)	1997/2000 (16.77 \pm 6.66)	1956/2000 (8.24 \pm 4.11)	273/2000 (13.7%)
	0.0	1894/2000 (5.61 \pm 3.12)	994/2000 (0.05 \pm 2.96)	1960/2000 (13.24 \pm 6.71)	1988/2000 (11.45 \pm 4.43)	494/2000 (24.7%)
	+0.4	1868/2000 (4.61 \pm 2.86)	869/2000 (-0.39 \pm 2.52)	1881/2000 (9.63 \pm 6.18)	1999/2000 (14.05 \pm 4.10)	801/2000 (40.1%)
	+0.8	1671/2000 (2.28 \pm 2.00)	921/2000 (-0.24 \pm 1.77)	1784/2000 (4.85 \pm 4.07)	2000/2000 (15.58 \pm 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 A8.4: 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
2×2	-0.8	389/2000 (1.49 \pm 3.84)	443/2000 (0.63 \pm 1.78)	1483/2000 (9.49 \pm 8.35)	1892/2000 (18.49 \pm 9.09)	271/2000 (-0.86 \pm 3.82)
	-0.4	672/2000 (2.58 \pm 4.59)	616/2000 (1.17 \pm 2.58)	1435/2000 (8.11 \pm 7.59)	1813/2000 (15.50 \pm 9.34)	361/2000 (-1.42 \pm 5.05)
	0.0	857/2000 (2.86 \pm 4.62)	746/2000 (1.48 \pm 2.96)	1413/2000 (6.76 \pm 6.85)	1727/2000 (12.72 \pm 9.29)	418/2000 (-1.38 \pm 5.36)
	+0.4	1009/2000 (3.05 \pm 4.38)	869/2000 (1.90 \pm 3.29)	1379/2000 (5.58 \pm 6.00)	1587/2000 (9.22 \pm 8.17)	510/2000 (-1.16 \pm 5.39)
	+0.8	1134/2000 (3.02 \pm 4.04)	1044/2000 (2.18 \pm 3.24)	1308/2000 (4.04 \pm 4.65)	1417/2000 (5.40 \pm 5.66)	638/2000 (-0.85 \pm 4.99)
	3×3	-0.8	1304/2000 (5.52 \pm 6.96)	1489/2000 (6.81 \pm 7.68)	1831/2000 (17.73 \pm 13.34)	2000/2000 (39.15 \pm 13.53)
-0.4		1410/2000 (5.63 \pm 6.43)	1630/2000 (7.10 \pm 6.85)	1810/2000 (15.14 \pm 11.72)	1998/2000 (33.65 \pm 12.92)	897/2000 (1.46 \pm 6.91)
0.0		1443/2000 (5.65 \pm 6.28)	1690/2000 (7.44 \pm 6.69)	1780/2000 (12.72 \pm 10.05)	1982/2000 (27.84 \pm 12.03)	1016/2000 (1.79 \pm 7.37)
+0.4		1511/2000 (5.55 \pm 5.79)	1756/2000 (7.35 \pm 6.07)	1766/2000 (10.27 \pm 8.43)	1957/2000 (20.49 \pm 10.99)	1119/2000 (1.80 \pm 7.17)
+0.8		1578/2000 (4.43 \pm 4.43)	1781/2000 (6.50 \pm 5.22)	1753/2000 (6.81 \pm 5.74)	1917/2000 (10.62 \pm 6.94)	1197/2000 (2.07 \pm 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 territory¹². 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).




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,

¹²Bolivia 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 [14](#), 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](#)).

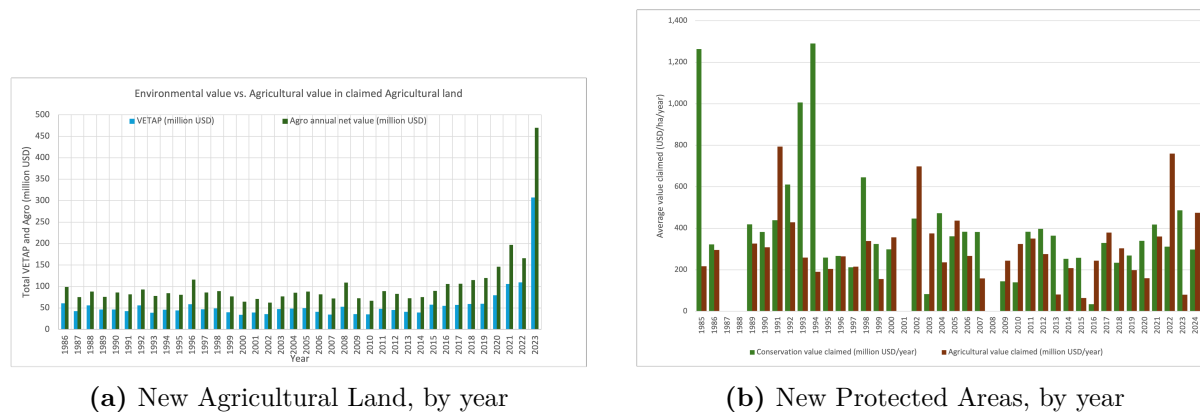
Following the economic downturn of the late 2010s and the subsequent pandemic crisis (2020–2023), Bolivia again intensified conservation efforts, rapidly designating new protected areas—particularly at the municipal level; the year 2019 saw the establishment of more than 4.2 million hectares of newly 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. \(2025\)](#). 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 berries¹³. 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 potential¹⁴.

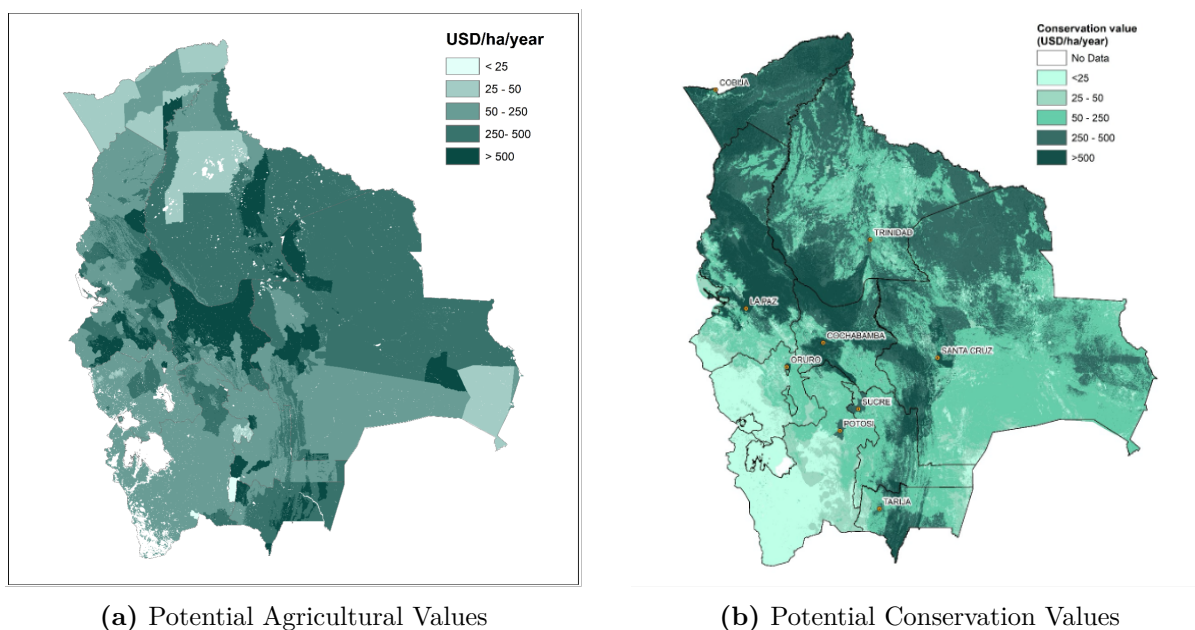


Figure A9.4: Potential Agricultural & Conservation values (annual USD per ha)

¹³In 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.

¹⁴While 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

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

denote the full sets of pixel-level values. Define the q th percentiles

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

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):

$$(46) \quad \mu_j^{(e)} = \frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} e_i, \quad \mu_j^{(a)} = \frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} a_i,$$

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:

$$(47) \quad f_j^{(e)}(q) = \frac{\#\{i \in \mathcal{I}_j : e_i \geq P_q^{(e)}\}}{|\mathcal{I}_j|}, \quad q \in \{90, 95, 99\},$$

and analogously for agricultural value:

$$(48) \quad f_j^{(a)}(q) = \frac{\#\{i \in \mathcal{I}_j : a_i \geq P_q^{(a)}\}}{|\mathcal{I}_j|}, \quad q \in \{90, 95, 99\}.$$

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:

$$(49) \quad \text{TAPS}_j^{(e)}(90) = 90 \cdot f_j^{(e)}(90), \quad \text{TAPS}_j^{(e)}(95) = 95 \cdot f_j^{(e)}(95), \quad \text{TAPS}_j^{(e)}(99) = 99 \cdot f_j^{(e)}(99).$$

and similarly for agricultural values:

$$(50) \quad \text{TAPS}_j^{(a)}(90) = 90 \cdot f_j^{(a)}(90), \quad \text{TAPS}_j^{(a)}(95) = 95 \cdot f_j^{(a)}(95), \quad \text{TAPS}_j^{(a)}(99) = 99 \cdot f_j^{(a)}(99).$$

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

$$(51) \quad M_j^{(e)} = F_\mu^{(e)}(\mu_j^{(e)}),$$

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:

$$(52) \quad T_j^{(e,90)}, T_j^{(e,95)}, T_j^{(e,99)}$$

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

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

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:

$$(53) \quad S_j^{(e)} = 0.40 M_j^{(e)} + 0.25 T_j^{(e,90)} + 0.20 T_j^{(e,95)} + 0.15 (T_j^{(e,99)})^{1.5},$$

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:

$$(54) \quad S_j^{(a)} = 0.40 M_j^{(a)} + 0.25 T_j^{(a,90)} + 0.20 T_j^{(a,95)} + 0.15 (T_j^{(a,99)})^{1.5}.$$

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: (DOI: 10.5281/zenodo.19613421)

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.