Online Supplementary Appendix to

"Adversarial Procurement in Two-Value Space: Cost-Effectiveness and the Knapsack Reversal"

Diana Weinhold   Lykke E. Andersen

August 2026

The main paper is available at

https://github.com/dmweinhold/Latest-Version-Papers/raw/main/Adversarial_Procurement.pdf

Appendix

This Online Supplementary Appendix collects the remaining formal material, equilibrium diagnostics, supplementary simulations, robustness exercises, Bolivia data-construction details, replication resources, and methodological disclosures for "Adversarial Procurement in Two-Value Space." The finite-board Claims-World priority-tilt theorem is proved in Appendix A of the paper.

The Supplementary Appendix is organised in three parts. Part I contains Budget-World formal material and equilibrium diagnostics. Part II reports supplementary simulation results and robustness exercises. Part III contains the Bolivia supplement, replication resources, and the AI-assisted methodology disclosure.

Contents

Part I. Budget-World formal material and equilibrium diagnostics

S1 Budget-World formal setup and theory

S1.1 Simulation setup

Economic interpretation and proportional-price normalisation

Let $a_i > 0$ denote the surplus generated by developing plot i, net of non-land development costs but before payment for the plot. Suppose the acquisition price capitalises a constant share $\theta \in (0,1)$ of that surplus:

$$c_i = \theta a_i, \qquad \pi_i = (1-\theta)a_i,$$

where π_i is Farmer's net payoff from acquiring and developing plot i. In economic units, initial budgets are

$$B_f^0 = s\sum_{i=1}^{N} c_i, \qquad B_g^0 = (1-s)\sum_{i=1}^{N} c_i.$$

Define normalized monetary quantities by

$$\widetilde{c}_i := \frac{c_i}{\theta} = a_i, \qquad \widetilde{B}_T := \frac{B_T}{\theta}, \quad T \in \{f, g\}.$$

Then, at every history,

$$c_i \le B_T \quad \Longleftrightarrow \quad a_i \le \widetilde{B}_T,$$

and a purchase update $B_T \leftarrow B_T - c_i$ is equivalent to $\widetilde{B}_T \leftarrow \widetilde{B}_T - a_i$. Likewise, the leakage update

$$B_f \leftarrow B_f - (1 - L)c_i$$

is equivalent to

$$\widetilde{B}_f \leftarrow \widetilde{B}_f - (1 - L)a_i.$$

For every priority tilt λ,

$$e_i c_i^{-\lambda} = \theta^{-\lambda} e_i a_i^{-\lambda},$$

so Green's priority order is unchanged. Farmer's terminal net payoff is $(1 - \theta) \sum_{i \in F} a_i$, a positive multiple of the objective $\sum_{i \in F} a_i$ used in the normalized game. Hence fixed-rule paths, best responses, and subgame-perfect equilibria correspond exactly between the economic and normalized representations. We therefore use normalized units throughout the formal analysis and write $\text{Price}(p_i) = a_i$; this is a normalization of the common proportionality coefficient, not an assumption of zero Farmer surplus.

Budgets

Let $\{(e_i, a_i)\}_{i=1}^N$ denote the environmental and agricultural values of $N = n^2$ plots. In the normalized units established above, $\text{Price}(p_i) = a_i$.

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. **Rival-Blind budgeter:** purchase the highest-a_i plots until budget is exhausted.

2. **Rival-Aware 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.

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 maximising $e_i \cdot a_i$.

3. **Block Farmers:** purchase the affordable plot with highest a_i.

4. **Max Efficiency (ratio-greedy):** purchase the affordable plot maximising $r_i = e_i/a_i$.

   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.[1]

Shorthand used throughout the proofs. ME always denotes **Max Environment** (the value-first

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

rule, ranking plots by e alone), and MX always denotes **Max Efficiency** (the ratio-greedy rule, ranking by e/a). The abbreviation ME stands for Max *Environment*, not efficiency; the efficiency rule is MX. Where they appear, Hot Spot (ranking by $e\,a$) and Block Farmers (ranking by a) are written HS and Block.

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. Throughout this appendix, Residual := $U^{\text{final}} = P \setminus (S_f^{\leq R} \cup S_g^{\leq R})$ denotes the set of plots not purchased by either team at the end of the game (using the notation of Section 5.1).

- **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.

S1.2 Budget World: constructive counterexamples and the case for asymptotic analysis

The exact finite Budget-World game with equal budgets does not support a blanket dominance theorem on Purchased Conservation: neither Max Environment nor Max Efficiency is universally better than the other across all finite instances. The proposition below establishes this through two constructive counterexamples. These examples motivate the asymptotic Budget-World analysis in Sections S1.5 and S1.8, which shows that the qualitative ranking does hold under the paper's baseline data-generating process as N grows. The simulation evidence on the baseline benchmark grids ($N = 100$) is reported in Section 5.3 of the main text and corroborates the asymptotic conclusion.

The constructions below 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 S1.1.

Proposition S1.1 (No universal finite-instance dominance in the equal-budget procurement game)**.** *Even with Rival-Blind 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. In particular, there exist instances on which Max Environment strictly outperforms Max Efficiency (Example (a)), and there exist instances on which Max Efficiency strictly outperforms Max Environment (Example (b)).*

Proof. Example (a): Max Environment beats Max Efficiency. Consider the six-plot board

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

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

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: $\text{PC} = 5+3+2 = 10$.

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: $\text{PC} = 10+5+3 = 18$.

Thus Max Environment strictly outperforms Max Efficiency on this instance.

Example (b): Max Efficiency beats Max Environment. 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 Rival-Blind 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 S1.2 (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.

S1.3 Asymptotic dominance: proof overviews

Shorthand used in the proofs. ME denotes **Max Environment**, the value-first rule ranking plots by E alone, and MX denotes **Max Efficiency**, the ratio-greedy rule ranking by E/A. Where they appear, Hot Spot (ranking by $E\,A$) and Block Farmers (ranking by A) are written HS and Block.

The finite counterexamples above show that neither ME nor MX dominates universally in arbitrary finite Budget-World instances. Theorem 4.1, stated in Section 5.3, proves a different claim: under the paper's baseline data-generating process, the knapsack reversal is a genuine asymptotic property. At $\rho = 0$, independence yields a closed-form fluid trajectory. Theorem 4.2 extends the realised ME $-$ MX gap to every $\rho \in [0, 0.3]$: a diagonal-shielding reduction identifies the gap, a uniform ratio-reservoir argument links that object to realised play, and verified interval arithmetic establishes strict correlation monotonicity. The earlier $\rho = 0.3$ frontier/value certificate is retained as an independent cross-check.

Both cases share one architecture, the *shielding identity* of Section S1.4. Writing $V_F(\rho) = \mathbb{E}_\rho[E\,\mathbf{1}\{A \geq q_F\}]$ for the static high-A block and $\mathcal{E}_F(\rho)$ for Farmer's realised ecological capture along the value-first path, ratio-greedy asymptotically buys the complementary low-A block

$V_G(\rho) = \mu_E - V_F(\rho)$, so the normalised gap is

$$\Delta(\rho) = V_F(\rho) - \mathcal{E}_F(\rho)$$

(Lemma S1.5): the ecological mass that value-first shields from the rival. The MX side $V_G(\rho)$ is closed-form and monotone in ρ (Lemmas S1.3 and S1.4). The remaining task is to identify Farmer's realised ecological capture $\mathcal{E}_F(\rho)$: independence yields a closed-form expression at $\rho = 0$, while the diagonal-collapse and cost-pinned-wedge arguments provide its static representation throughout $[0, 0.3]$. The uniform ratio-reservoir certificate then identifies the resulting shielding gap with the realised ME − MX gap over that entire interval.

The detailed baseline proof is in Section S1.5, with supporting finite-N inputs collected in Section S1.6. The positive-correlation extension is in Section S1.8.

Proof overview for Theorem 4.1. At $\rho = 0$, the baseline law has independent identical E and A marginals. In the fluid limit, the tail-mass coordinates of the Farmer and Green frontiers move symmetrically:

$$u(s) = v(s) = 1 - \sqrt{1 - 2s}.$$

This symmetry makes Farmer's ecological capture along the Farmer-active ME path equal Green's agricultural spend, so the dynamic object is itself closed-form:

$$\mathcal{E}_F(0) = \lim_N \frac{1}{N} \sum_{t \le \tau N} E(f_t^{\mathrm{ME}}) = S_G(\tau), \qquad \text{whence} \qquad \frac{1}{N}\mathrm{PC}_N^{\mathrm{ME}} \longrightarrow \mu_E - S_G(\tau).$$

On the MX side, the safe-frontier argument shows ratio-greedy buys the static block $\{A \le q_F\}$, so by independence—the $\rho = 0$ case of Lemma S1.3—

$$\frac{1}{N}\mathrm{PC}_N^{\mathrm{MX}} \longrightarrow V_G(0) = \mu_E(1 - \phi), \qquad V_F(0) = \mu_E \phi.$$

The shielding identity then gives the exact asymptotic gap in closed form,

$$\Delta(0) = V_F(0) - \mathcal{E}_F(0) = \mu_E\phi - S_G(\tau) = 0.216135\ldots > 0.$$

The proof in Section S1.5 records the closed-form trajectory, the symmetry lemma, and the interval-certified numerical margin.

Proof overview for Theorem 4.2. At $\rho = 0.3$ the independence symmetry disappears, but the gap admits an exact static reduction (Section S1.8.1). The proof works in latent Gaussian coordinates $X := Z_A$, $Y := Z_E$ and proceeds in five steps.

First, the static MX side reduces to a closed form: the safe-frontier argument shows ratio-greedy buys the low-A block up to $o_p(N)$ boundary labels, and Lemma S1.3 evaluates that block exactly,

$$\frac{1}{N}\mathrm{PC}_N^{\mathrm{MX}} \longrightarrow V_G(0.3) = 3.265118\ldots.$$

Second, on the ME side the value-first active frontier collapses onto the diagonal $Y = X$ (Lemma S1.27): the gap between the two latent thresholds contracts, so the frontier is governed by a single scalar cutoff.

Third, Farmer's realised value-first capture becomes a static, cost-pinned *diagonal wedge* (Lemma S1.28): $\mathcal{E}_F(0.3)$ is the E-content of $F_\rho = \{X \geq x_\rho,\ Y \leq X\}$, with the cutoff x_ρ fixed by the budget identity $\mathbb{E}[\mathsf{m}(X)\mathbf{1}_{F_\rho}] = \frac{1}{2}\mu_A$.

Fourth, the gap $\Delta(\rho) = V_F(\rho) - \mathcal{E}_F(\rho)$ is a *cost-equal exchange* across the diagonal (Lemma S1.29), nonnegative because $Y = X$ is exactly the $E = A$ contour.

Fifth, certification is static rather than dynamic. Interval arithmetic encloses the exact value $\Delta(0.3)$ (Theorem S1.30), and a 60-box derivative sweep certifies $\Delta'(\rho) > 0.20$ on

$[0, 0.3]$ (Lemma S1.31), so Δ is strictly increasing there with $\Delta(0.3) = 0.293565\ldots > \Delta(0) = 0.216135\ldots$. Identifying Δ with the realised $\mathrm{ME} - \mathrm{MX}$ gap uses the MX static-block limit, established uniformly on $[0, 0.3]$ by the ratio-reservoir argument of Appendix S1.7.

The earlier frontier/value certificate—an explicit fluid-limit ODE integrated with a validated Picard step-tube and a stopping-time affordability bootstrap—is retained as an independent cross-check (Section S1.8). It certifies the same $\rho = 0.3$ dominance by a separate route, $\liminf_{N\to\infty} \frac{1}{N}\mathbb{E}[\mathrm{PC}_N^{\mathrm{ME}} - \mathrm{PC}_N^{\mathrm{MX}}] \geq 0.198653\ldots > 0$, and supplies the MX threshold identity and entrance-uniqueness lemma that the diagonal-shielding argument cites. The two routes agree, the exact static value $0.293565\ldots$ strengthening the earlier lower bound.

Scope and limitations. The Budget-World theorems are formal asymptotic results for i.i.d. boards with uniform marginals on $[0.1, 10]$ coupled by a Gaussian copula with latent correlation $\rho \in [0, 0.3]$. Theorem 4.1 gives the closed-form independence benchmark, $\Delta(0) = 0.216135\ldots > 0$. Theorem 4.2 extends the result to the full interval: for every $\rho \in [0, 0.3]$, the normalised expected Purchased Conservation gap

$$\Delta(\rho) = \lim_{N\to\infty} \frac{1}{N}\mathbb{E}\Big[\mathrm{PC}_N^{\mathrm{ME}} - \mathrm{PC}_N^{\mathrm{MX}}\Big]$$

exists, coincides with the diagonal-shielding gap, and is strictly increasing, with

$$\Delta(0.3) = 0.293565\ldots > \Delta(0) = 0.216135\ldots > 0.$$

Thus the certified result is an interval theorem, not merely a pair of endpoint results. Simulations suggest that the normalised gap remains positive beyond the proved range and follows a hump-shaped profile in ρ. The ratio-reservoir argument that identifies the shielding gap with the realised $\mathrm{ME} - \mathrm{MX}$ gap is, however, certified only on $[0, 0.3]$, so the present theorem does not cover the high-correlation regime. The results are also asymptotic: the finite counterexamples above show that small or adversarially constructed boards can reverse the comparison.

The Budget-World proof chain is specific to the paper's benchmark mechanism: the proportional-price specification, normalised to $c_i = A_i$, parity budgets, full leakage, Farmer moving first, and the Rival-Blind max-A Farmer rule. It compares value-first (ME) and ratio-greedy (MX) on Purchased Conservation; it does not establish finite-board dominance or a theorem for strategic Farmer play, unequal budget shares, incomplete leakage, non-proportional price links, alternative marginal laws, or correlations outside $[0, 0.3]$. The main paper's Claims-World theorems provide separate, distribution-free finite-board results for the free-claiming game. The small-N SPE exercises and the ecological selection premium should be read as scope diagnostics: they help identify when the reversal is likely to be strong, but they are not additional formal Budget-World dominance results.

S1.4 Static blocks, closed form, and the shielding identity

Both asymptotic Budget-World theorems compare value-first against ratio-greedy through the same pair of static *blocks*, split at the Farmer A-frontier. We collect the shared notation and three distribution-level facts here, then instantiate them at $\rho = 0$ (Section S1.5) and $\rho = 0.3$ (Section S1.8). Throughout, $\ell = 0.1$, $\nu = 10$, $\mu_E = \mu_A = (\ell + \nu)/2 = 5.05$, and the marginals are uniform on $[\ell, \nu]$ under a Gaussian copula with latent correlation ρ.

The Farmer frontier $q_F = \sqrt{50.005}$ is fixed by the half-budget condition $\int_{q_F}^{\nu} a \, \frac{da}{\nu - \ell} = \frac{1}{2}\mu_A$ and depends only on the A-marginal; set

$$\phi := \Pr(A \geq q_F) = \frac{\nu - q_F}{\nu - \ell} = 0.295816, \qquad z_F := \Phi^{-1}(1 - \phi) = 0.536472,$$

both ρ-invariant. Define the *high-A* and *low-A block values*

$$V_F(\rho) := \mathbb{E}_\rho[E\,\mathbf{1}\{A \geq q_F\}], \qquad V_G(\rho) := \mathbb{E}_\rho[E\,\mathbf{1}\{A \leq q_F\}], \qquad V_F(\rho) + V_G(\rho) = \mu_E.$$

Lemma S1.3 (Closed form for the block values)**.** *For every $\rho \in (-1,1)$,*

$$V_G(\rho) = \ell\,(1-\phi) + (\nu-\ell)\,\Phi_2\Big(0,\, z_F;\, -\tfrac{\rho}{\sqrt{2}}\Big), \qquad V_F(\rho) = \mu_E - V_G(\rho), \tag{S1.1}$$

where $\Phi_2(\cdot,\cdot;r)$ is the standard bivariate-normal CDF with correlation r. At $\rho = 0$ the copula factorises and $V_G(0) = \mu_E(1-\phi)$, $V_F(0) = \mu_E\phi = 1.493871\ldots$; at $\rho = 0.3$, $V_G(0.3) = 3.265118\ldots$ and $V_F(0.3) = 1.784882\ldots$.

Proof. Write $E = \ell + (\nu-\ell)\Phi(Z_E)$ and $\{A \leq q_F\} = \{Z_A \leq z_F\}$, so $V_G(\rho) = \ell(1-\phi) + (\nu - \ell)\,\mathbb{E}_\rho[\Phi(Z_E)\mathbf{1}\{Z_A \leq z_F\}]$. Introduce $W \sim N(0,1)$ independent of (Z_E, Z_A); since $\Phi(Z_E) = \Pr(W \leq Z_E \mid Z_E)$,

$$\mathbb{E}_\rho[\Phi(Z_E)\mathbf{1}\{Z_A \leq z_F\}] = \Pr(W - Z_E \leq 0,\ Z_A \leq z_F).$$

The pair $(W - Z_E, Z_A)$ is centred bivariate normal with $\mathrm{Var}(W - Z_E) = 2$, $\mathrm{Var}(Z_A) = 1$, and $\mathrm{Cov}(W - Z_E, Z_A) = -\rho$, hence correlation $-\rho/\sqrt{2}$; standardising the first coordinate gives $\Phi_2(0, z_F; -\rho/\sqrt{2})$, which is (S1.1). At $\rho = 0$, $\Phi_2(0, z_F; 0) = \Phi(0)\Phi(z_F) = \frac{1}{2}(1-\phi)$, so $V_G(0) = (1-\phi)(\ell + \frac{\nu-\ell}{2}) = \mu_E(1-\phi)$. The numerical values are single evaluations of the bivariate-normal CDF (Genz, 2004); the $\rho = 0.3$ value lies in the archived enclosure $[3.26511, 3.26513]$. $\square$

Lemma S1.4 (Correlation monotonicity of the block values)**.** *$V_G(\rho)$ is non-increasing and $V_F(\rho)$ non-decreasing in ρ, for every $\rho \in (-1,1)$.*

Proof. By (S1.1) it suffices that $r \mapsto \Phi_2(0, z_F; r)$ is non-decreasing, since $r = -\rho/\sqrt{2}$ decreases in ρ. Plackett's identity Plackett (1954) gives $\partial_r \Phi_2(a, b; r) = \varphi_2(a, b; r) \geq 0$, so $\Phi_2(0, z_F; \cdot)$ is non-decreasing and V_G is non-increasing. As q_F is ρ-invariant, $V_F = \mu_E - V_G$ inverts the sign. $\square$

Lemma S1.5 (Shielding identity)**.** *Suppose the ratio-greedy limit* $N^{-1}\mathrm{PC}_N^{\mathrm{MX}} \to V_G(\rho)$ *holds. Along any subsequence* (N_k) *on which* $N_k^{-1}\mathrm{PC}_{N_k}^{\mathrm{ME}}$ *converges, write* $\mathcal{E}_F^\star(\rho) := \mu_E - \lim_k N_k^{-1}\mathrm{PC}_{N_k}^{\mathrm{ME}}$ *for the corresponding Farmer ecological capture along the value-first path. Then the subsequential payoff gap satisfies*

$$\lim_k N_k^{-1}(\mathrm{PC}_{N_k}^{\mathrm{ME}} - \mathrm{PC}_{N_k}^{\mathrm{MX}}) = (\mu_E - \mathcal{E}_F^\star(\rho)) - V_G(\rho) = V_F(\rho) - \mathcal{E}_F^\star(\rho).$$

In particular, if the full ME *limit exists, then with* $\mathcal{E}_F(\rho) := \mu_E - \lim_N N^{-1}\mathrm{PC}_N^{\mathrm{ME}}$*,*

$$\Delta(\rho) := \lim_N \tfrac{1}{N}\,\mathbb{E}[\mathrm{PC}_N^{\mathrm{ME}} - \mathrm{PC}_N^{\mathrm{MX}}] = V_F(\rho) - \mathcal{E}_F(\rho).$$

Either way the gap is the ecological mass of the high-A block that value-first shields from Farmer: the MX *side is closed-form (Lemma S1.3) and monotone (Lemma S1.4), so only the Farmer capture carries dynamic frontier-path content.*

Proof. Immediate from $V_F + V_G = \mu_E$ and the definitions of the Farmer captures. $\square$

S1.5 Baseline asymptotic dominance theorem

This subsection proves the Budget-World theorem under the independent baseline law:

$$(E_i, A_i)_{i=1}^N \overset{iid}{\sim} U[0.1, 10] \times U[0.1, 10], \qquad E \perp A,$$

with parity budgets, Farmer moving first, full leakage, Purchased Conservation as the outcome, ME = affordable max-E, and MX = affordable max-E/A. Write $\ell = 0.1$, $\nu = 10$, and $\mu_E = \mu_A = 5.05$.

S1.5.1 Fluid trajectory and the symmetry identity

Let $u(s) = 1 - F_U(x(s))$ and $v(s) = 1 - F_U(c(s))$ be the tail masses above the current Farmer and Green frontiers. The survivor-law fluid limit gives

$$\dot{u} = \frac{1}{1-v}, \qquad \dot{v} = \frac{1}{1-u}, \qquad u(0) = v(0) = 0.$$

By symmetry,

$$u(s) = v(s) = 1 - \sqrt{1-2s}, \qquad s \in [0, 1/2]. \tag{S1.2}$$

Substitution into the budget equations gives

$$S_F(s) = 0.1s + 3.3\{1 - (1-2s)^{3/2}\}, \tag{S1.3}$$

$$S_G(s) = 0.1s + 1.65\{1 - (1-2s)^{3/2}\}. \tag{S1.4}$$

Let τ solve $S_F(\tau) = 2.525$. Interval evaluation gives

$$0.304718 < \tau < 0.304719.$$

After Farmer exhaustion, the Green-only phase shops from the frozen frontier and the post-Farmer budget identity gives

$$\kappa = 1 - \tau. \tag{S1.5}$$

The finite-game endpoint error is $O(1)$ labels by Proposition S1.13.

Lemma S1.6 (E–A symmetry of per-turn captures)**.** *For $s \in [0, \tau]$, conditional on the fluid state at share s, Farmer's per-turn ecological capture and Green's per-turn agricultural spend have the same limit:*

$$\mathbb{E}[E(f^N_{\lfloor sN \rfloor})] = \mathbb{E}[A(g^N_{\lfloor sN \rfloor})] + o(1) = \frac{\ell + x(s)}{2} + o(1).$$

Consequently,

$$\frac{1}{N} \sum_{t=1}^{\lfloor \tau N \rfloor} E(f^N_t) \xrightarrow{p} S_G(\tau). \tag{S1.6}$$

Proof. Immediately before the sN-th Farmer move, the survivor law is the doubly truncated host $\nu_{x(s),c(s)}$. Under independence this host factors: $A \sim U[\ell, x(s)]$ and $E \sim U[\ell, c(s)]$. Farmer's max-A pick has an E-coordinate distributed as the truncated E-marginal, while Green's max-E pick has an A-coordinate distributed as the truncated A-marginal. Since $x(s) = c(s)$ by (S1.2), the two conditional expectations agree. The integral over $s \in [0, \tau]$ is exactly the Green agricultural spend path $S_G(\tau)$. $\square$

Lemma S1.7 (Bounded martingale aggregation for the baseline path)**.** *Let $(\mathcal{F}_t)$ be the natural filtration immediately before Farmer's t-th move on the* ME *path, and let $T_N \le CN$ be any stopping time. For bounded turn values $Z_t \in [\ell, \nu]$,*

$$M_N(T_N) := \frac{1}{N} \sum_{t=1}^{T_N} \{Z_t - \mathbb{E}[Z_t \mid \mathcal{F}_t]\}$$

satisfies $M_N(T_N) \to 0$ in L^2, hence in probability. In particular, if

$$\sup_{0 \le s \le T} \left| \mathbb{E}[Z_{\lfloor sN \rfloor} \mid \mathcal{F}_{\lfloor sN \rfloor}] - z(s) \right| \xrightarrow{p} 0$$

for a bounded continuous z, then

$$\frac{1}{N}\sum_{t=1}^{\lfloor TN\rfloor} Z_t \xrightarrow{p} \int_0^T z(s)\,ds.$$

The same conclusion holds with $\lfloor TN\rfloor$ replaced by a stopping time T_N satisfying $T_N/N \to T$ in probability.

Proof. The increments

$$D_t := Z_t - \mathbb{E}[Z_t \mid \mathcal{F}_t]$$

are martingale differences and satisfy $|D_t| \le \nu - \ell$. Therefore

$$\mathbb{E}[M_N(T_N)^2] \le \frac{1}{N^2}\mathbb{E}\left[\sum_{t\le T_N} D_t^2\right] \le \frac{C(\nu-\ell)^2}{N} \to 0.$$

The predictable compensator converges to the integral by the assumed uniform drift convergence and Riemann-sum convergence. The random-time replacement uses boundedness:

$$\left|\frac{1}{N}\sum_{t\le T_N} Z_t - \frac{1}{N}\sum_{t\le\lfloor TN\rfloor} Z_t\right| \le \nu\,\frac{|T_N - \lfloor TN\rfloor|}{N} = o_p(1).$$

□

Theorem S1.8 (Baseline asymptotic dominance theorem, exact form)**.** *Under the baseline Budget-World theorem object,*

$$\lim_{N\to\infty}\frac{1}{N}\mathbb{E}[\mathrm{PC}_N^{\mathrm{ME}} - \mathrm{PC}_N^{\mathrm{MX}}] = \mu_E\phi - S_G(\tau) = 0.216135\ldots > 0,$$

where ϕ is the MX *Farmer-exhaustion share and τ is the* ME *Farmer-exhaustion share.*

Proof. For ME, the support inputs imply exact max-E play through the Farmer-active window and leave only $o_p(N)$ labels unresolved at the endpoint: early affordability is Lemma S1.10, compact affordability is Lemma S1.11, the Farmer crossing is Proposition S1.12, the leftover bound is Proposition S1.13, and deterministic replacement of the random stopping index is Lemma S1.16. Hence

$$\frac{1}{N}\mathrm{PC}_N^{\mathrm{ME}} = \frac{1}{N}\sum_{i=1}^{N} E_i - \frac{1}{N}\sum_{t=1}^{K_N^F(\mathrm{ME})} E(f_t^N) + o_p(1) \xrightarrow{p} \mu_E - S_G(\tau),$$

where Lemmas S1.6 and S1.7 identify the predictable Farmer-ecology compensator with $S_G(\tau)$ and remove the martingale noise.

For MX, Proposition S1.20 gives $N^{-1}\mathrm{PC}_N^{\mathrm{MX}} \to V_G(0)$, the low-$A$ block value, which by independence is the $\rho = 0$ case of the closed form (Lemma S1.3):

$$\frac{1}{N}\mathrm{PC}_N^{\mathrm{MX}} \xrightarrow{p} V_G(0) = \mathbb{E}[E\,\mathbf{1}\{A \le q_F\}] = \mu_E(1-\phi).$$

Both normalised payoffs are bounded by 10, so convergence in probability implies convergence of expectations. Writing $\mathcal{E}_F(0) := \mu_E - \lim_N N^{-1}\mathrm{PC}_N^{\mathrm{ME}} = S_G(\tau)$ (the ME limit above) and $V_F(0) = \mu_E\phi$, the shielding identity (Lemma S1.5) gives

$$\lim_{N\to\infty} \frac{1}{N}\mathbb{E}[\mathrm{PC}_N^{\mathrm{ME}} - \mathrm{PC}_N^{\mathrm{MX}}] = V_F(0) - \mathcal{E}_F(0) = \mu_E\phi - S_G(\tau).$$

Finally,

$$\phi = \frac{10 - \sqrt{50.005}}{9.9} = 0.295816\ldots, \qquad \mu_E\phi = 1.493871\ldots,$$

and the interval $0.304718 < \tau < 0.304719$ gives

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

Thus the certified asymptotic gap lies in

$$0.216133 < \mu_E \phi - S_G(\tau) < 0.216137.$$

□

Remark S1.9 (Interpretation and scope). At $\rho = 0$ the shielding identity (Lemma S1.5) is closed-form on both sides: independence gives $V_F(0) = \mu_E \phi$, and the symmetry identity gives $\mathcal{E}_F(0) = S_G(\tau)$, so $\Delta(0) = \mu_E \phi - S_G(\tau)$. MX cedes a ϕ-fraction of the market, worth $\mu_E \phi$ ecologically under independence, while value-first keeps all but $S_G(\tau)$. Both simplifications are special to independent identical marginals; at $\rho = 0.3$ (Appendix S1.8) the block value V_F stays closed-form via Lemma S1.3 while F requires the frontier/value-path certificate.

S1.6 Baseline proof tools

Frozen object and constants

We keep the notation from Appendix S1.5. Write

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

The ratio score for MX is

$$R := E/A.$$

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

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

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

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

Under the baseline uniform law,

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

The baseline proof uses four standard inputs: thresholded empirical convergence, deterministic replacement of convergent stopping times, safe-frontier exactness for the ratio rule through Farmer exhaustion, and a one-dimensional density-dependent LLN for block depletion.

Lemma S1.10 (Early baseline affordability)**.** *Let $a < \mu_A/(2\overline{u})$, where $\overline{u} = 10$ 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, affordability is inactive on the prefix $[0, m]$ for $m = 0.10$.*

Proof. By the law of large numbers, $B_G^0/N, B_F^0/N \xrightarrow{p} \mu_A/2$. Through $\lfloor aN \rfloor$ own moves, either player can spend at most $\overline{u}aN + O(1)$. Since $\overline{u}a < \mu_A/2$, choose $\eta > 0$ with $\overline{u}a + 2\eta < \mu_A/2$. With probability tending to one, each initial budget is at least $(\mu_A/2 - \eta)N$, so after $\lfloor aN \rfloor$ own moves each residual budget is at least $\eta N + O(1)$. For all large N, this exceeds $\overline{u}$, hence every remaining plot is affordable. $\square$

Lemma S1.11 (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* (S1.2)*. Define*

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

Let

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

Then $\beta_\varepsilon > 0$*, and*

$$\mathbb{P}(\textit{Green plays exact max-}E \textit{ throughout } [m, \kappa - \varepsilon]) \to 1.$$

On this event the normalised frontier-budget process converges uniformly to its fluid limit on $[m, \kappa - \varepsilon]$*.*

Proof. The displayed path is the candidate Green budget under exact play. The first line is the Farmer-active spend path S_G. After Farmer exhaustion, the Farmer frontier is frozen at x_τ; under independence, Green's max-E pick has an A-coordinate drawn from the residual uniform marginal on $[\ell, x_\tau]$ up to an endpoint error, giving post-Farmer spending rate $(\ell + x_\tau)/2$. The two pieces concatenate continuously at τ, are strictly decreasing, and reach zero at $\kappa = 1 - \tau$. Hence $\beta_\varepsilon = b_G(\kappa - \varepsilon) > 0$.

Let σ_N be the first share $s \ge m$ at which some surviving plot has price larger than the absolute Green budget $N b_G^N(s)$, with $\sigma_N = \infty$ if no such failure occurs before κ. On the stopped interval before σ_N, Green plays exact max-E. The active phase and the post-Farmer Green-only phase then have the fluid drifts used above, with bounded normalised budget increments;

applying the standard density-dependent LLN on each phase and concatenating at τ gives

$$\sup_{m\le s\le(\kappa-\varepsilon)\wedge\sigma_N}\|Z_N(s)-Z(s)\|\xrightarrow{p}0,$$

where Z_N records the frontiers and normalised budgets.

If $\mathbb{P}(\sigma_N\le\kappa-\varepsilon)\not\to 0$, then along a subsequence on that event affordability failure implies $b_G^N(\sigma_N)<\overline{u}/N$. Stopped convergence instead implies

$$b_G^N(\sigma_N)\to b_G(\sigma_N\wedge(\kappa-\varepsilon))\ge\beta_\varepsilon>0$$

in probability, a contradiction. Thus $\mathbb{P}(\sigma_N>\kappa-\varepsilon)\to 1$, and the stopped convergence is unstopped on the claimed interval. $\square$

Proposition S1.12 (Baseline Farmer-exhaustion crossing time)**.** *Let τ be the unique solution of $S_F(\tau)=\mu_A/2$ on $(0,1/2)$, and let*

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

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

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

Proof. The random initial Farmer budget satisfies $B_F^0/N\to\mu_A/2$ 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

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

Uniform convergence $b_F^N \to b_F$ 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,

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

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

Proposition S1.13 (Baseline leftover-count bound at terminal time)**.** *Let $T^\star$ 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 $T^\star$ satisfies

$$|U_{T^\star}| \;\le\; \left\lceil \frac{2\nu}{\ell} \right\rceil \;=\; 200.$$

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

Proof. Parity budgets give the pathwise identity

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

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 $T^\star$, neither player can afford any remaining plot; in particular, $B_F^{\mathrm{res}}(T^\star) < \nu$ and $B_G^{\mathrm{res}}(T^\star) < \nu$. Therefore

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

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

For the consequence, Proposition S1.12 gives $K_N^F(\mathrm{ME})/N \xrightarrow{p} \tau$, and Lemma S1.11 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 $T^\star$ coincides with Green's stopping time, and Farmer is already unable to afford any remaining plot at $T^\star$. The leftover bound then reads

$$N - K_N^{\mathrm{ME}} - K_N^F(\mathrm{ME}) \;=\; |U_{T^\star}| \;\le\; 200$$

on that event. Dividing by N and combining with $K_N^F(\mathrm{ME})/N \xrightarrow{p} \tau$ yields $K_N^{\mathrm{ME}}/N \xrightarrow{p} 1-\tau$. $\square$

Tool 1. 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 S1.14 (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, and assume F_Z is strictly increasing on an open interval containing $Q_Z(K)$ (equivalently, Q_Z is continuous on K); under the continuous, positive-density laws used in this paper this holds. Then, for every compact $K \subset (0,1)$,*

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

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 (van der Vaart, 1998, Chapter 21) or Shorack and Wellner (Shorack and Wellner, 1986, Chapter 2). The indicator classes

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

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 $\hat{r}_{s,N} \to r_s$. □

Remark S1.15 (Where Tool 1 is used). S1 is the justification behind all static portfolio functionals in the baseline chain. In particular it identifies

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

as the large-N value and cost of the initial top-sN ratio portfolio.

Tool 2. 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 normalisation by N makes this harmless whenever split times converge and per-turn terms are bounded.

Lemma S1.16 (Deterministic replacement of convergent split times)**.** *Let τ_N be an integer-valued random time with $\tau_N/N \to \tau \in [0,1]$ in probability. Let $(Y_{N,t})_{1\le t\le N}$ be random variables with $|Y_{N,t}| \le M$ almost surely for all N, t. Then*

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

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

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

Proof. Pathwise,

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

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

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

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

Remark S1.17 (Where Tool 2 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 normalised value or cost sums. This is not deep; it just needs to be stated instead of waved away.

Tool 3. Safe-frontier exactness through Farmer exhaustion

This tool proves that, under the independent-uniform baseline, MX's Farmer-active purchases are asymptotically the static descending-ratio list and that Farmer exhausts at the benchmark share ϕ. The proof has two margins: the static top-ratio block is below the Farmer A-frontier, and its cumulative cost is far below Green's initial budget.

Under baseline independence and uniform marginals, for $1 \leq r \leq \nu/\ell$,

$$\mathbb{P}(R \geq r) = \frac{(\nu - \ell r)^2}{2r(\nu - \ell)^2}. \tag{S1.7}$$

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

$$r_\phi \approx 1.66752, \qquad q_F := Q_A(1 - \phi) \approx 7.07142, \qquad \frac{\nu}{r_\phi} \approx 5.99693 < q_F.$$

Since $s \mapsto \nu/r_s$ is increasing and $s \mapsto Q_A(1-s)$ is decreasing, this strict inequality is the worst case on the Farmer-active window.

Lemma S1.18 (Baseline ratio-side cost and safe-frontier margins). *Under the baseline independent-uniform law,*

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

is continuous and strictly decreasing. Moreover $C_R(r_\phi) \approx 0.6110 < \mu_A/2 = 2.525$ *and* $\nu/r_\phi <$

q_F. Hence there exist constants $\delta > 0$, $\eta > 0$, and $\varepsilon_0 > 0$ such that

$$C_R(r_{\phi+\delta}) \le \frac{\mu_A}{2} - \varepsilon_0, \qquad \frac{\nu}{r_{\phi+\delta} - 2\eta} < Q_A(1 - \phi - \delta) - 2\eta. \tag{S1.8}$$

Consequently, on the event that the empirical ratio and A-cutoffs at share $\phi + \delta$ are within η of their population values, the initial top-$(\phi + \delta)N$ ratio block is disjoint from the first $\lfloor(\phi + \delta)N\rfloor$ items of the descending A-order.

Proof. For $R = E/A$, the event $\{R \ge r\}$ is $\{(a, e) : e \ge ra\}$, hence $a \le \nu/r$. Integrating the joint uniform density on $[\ell, \nu]^2$ gives the displayed formula for C_R, and Leibniz differentiation gives

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

The numerical inequalities above have large strict margins. Continuity of C_R, $s \mapsto r_s$, and $s \mapsto Q_A(1 - s)$ gives (S1.8) for small δ, η. The empirical cutoff event has probability tending to one by Proposition S1.14. On that event, every item in the initial top-$(\phi + \delta)N$ ratio block has $A \le \nu/(r_{\phi+\delta} - \eta) < Q_A(1 - \phi - \delta) - \eta$, while every item in the first $\lfloor(\phi + \delta)N\rfloor$ descending-A block has $A \ge Q_A(1 - \phi - \delta) - \eta$. The two blocks are therefore disjoint. □

Corollary S1.19 (Green is unconstrained through the MX safe-frontier window, baseline)**.** *Let $\delta, \eta, \varepsilon_0$ be as in Lemma S1.18. With probability tending to one,* MX *Green plays exact* $\arg\max R$ *over the whole survivor set through round $\lfloor(\phi + \delta)N\rfloor$, not merely over the affordable subset.*

Proof. Let τ_N^G be Green's first affordability failure before $\lfloor(\phi+\delta)N\rfloor$, and work on the empirical safe-frontier event of Lemma S1.18. Up to τ_N^G, Green's purchases are the static descending-ratio items and, by disjointness, Farmer's purchases do not remove them. Hence Green's cumulative spend through any such round is bounded by the empirical cost of the initial top-$(\phi+\delta)N$ ratio block. By Proposition S1.14 and (S1.8), this cost is at most $(\mu_A/2 - \varepsilon_0/2)N$ with probability tending to one, while $B_G^0/N \to \mu_A/2$. Thus Green's residual budget is at least $\varepsilon_0 N/3$ throughout the window with probability tending to one. For large N, this exceeds ν, so every surviving plot is affordable and τ_N^G cannot occur. □

Proposition S1.20 (MX static-path exactness and Farmer exhaustion)**.** *Under the baseline law, the following hold.*

(i) For every deterministic $T < \phi$, with probability tending to one, Farmer's first $\lfloor TN \rfloor$ purchases under the MX *path are the initial descending-A items, while Green's first $\lfloor TN \rfloor$ purchases are the initial descending-ratio items.*

(ii) Consequently,

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

and the same identities hold at $T = \phi$ by compact-to-endpoint passage.

(iii) $K_N^F(\mathrm{MX})/N \xrightarrow{p} \phi$. Moreover the total MX *Green portfolio converges, up to $o_p(N)$ boundary and terminal labels, to the static region $\{A \le q_F\}$. Hence*

$$\frac{1}{N}\mathrm{PC}_N^{\mathrm{MX}} \xrightarrow{p} \mathbb{E}[E\mathbf{1}\{A \le q_F\}].$$

Proof. Let σ_N be the no-Green descending-A budget-crossing time,

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

where $A_{(1)} \ge \cdots \ge A_{(N)}$. By Proposition S1.14, the law of large numbers for B_F^0, and the defining identity $\int_0^\phi Q_A(1-u)du = \mu_A/2$, the standard monotone crossing argument gives $\sigma_N/N \xrightarrow{p} \phi$.

On the intersection of the safe-frontier event, the Green-affordability event from Corollary S1.19, and $\{\sigma_N \le (\phi+\delta)N\}$, induction gives the exact static path through σ_N: Farmer can afford and buys the next descending-A item, while Green buys the next descending-ratio item; the two blocks are disjoint, so neither player removes the other's next static item. This proves (i) for every fixed $T < \phi$. The convergence of normalised sums follows from Proposition S1.14. Since $E, A \in [\ell, \nu]$, the endpoint $T = \phi$ follows by bounding the compact-to-endpoint tail by $\nu(\phi - T) + o(1)$ and then taking $T \uparrow \phi$, proving (ii).

After σ_N, Farmer's residual budget is below ν. Since each remaining item costs at least ℓ, Farmer can make at most $\lceil \nu/\ell \rceil = 100$ further purchases. Hence $0 \le K_N^F(\mathrm{MX}) - \sigma_N \le 100$ with probability tending to one, proving $K_N^F(\mathrm{MX})/N \to \phi$. Farmer's active purchases are therefore the empirical top-A block $\{A > q_F\}$ up to $o_p(N)$ labels. Green's active ratio block lies inside $\{A \le q_F\}$ by the safe-frontier inequality, and after Farmer exhaustion Green purchases all but $O(1)$ of the remaining labels by Proposition S1.13. Thus Green's total MX portfolio is $\{A \le q_F\}$ up to $o_p(N)$ labels. The thresholded LLN gives the claimed normalised PC limit. □

Tool 4. Density-dependent LLN for the block-depletion chain

We use one standard fluid-limit input. If a scaled Markov chain has jumps of order $1/N$, uniformly convergent drift, a locally Lipschitz limiting vector field, and a convergent initial condition, then it converges uniformly on compact time intervals to the corresponding ODE solution (Kurtz, 1970; Ethier and Kurtz, 1986; Wormald, 1999).

Theorem S1.21 (Specialised density-dependent LLN)**.** *Let (Z_t^N) be a Markov chain on a compact interval and set $z_N(s) = Z^N_{\lfloor sN \rfloor}$. Suppose that, uniformly on compact subsets of the region visited by the limit path,*

$$|Z^N_{t+1} - Z^N_t| \le C/N, \qquad \mathbb{E}[Z^N_{t+1} - Z^N_t \mid \mathcal{F}_t] = N^{-1} f(Z^N_t, t/N) + o(N^{-1}),$$

where f is locally Lipschitz, and suppose $Z_0^N \to z_0$ in probability. Then the ODE $\dot{z} = f(z,s)$, $z(0) = z_0$, has a unique solution up to the horizon considered and

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

Proposition S1.22 (Verification for the baseline block-depletion chain)**.** *Fix $T < 1/2$ and $\alpha \in (0,1)$. Let $X_t^{N,\alpha}$ be the number of surviving labels in the fixed initial αN-block before Green turn t, and set $x_t^{N,\alpha} = X_t^{N,\alpha}/N$. On any compact interval on which the limiting count is strictly positive,*

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

with

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

Consequently $x_N^\alpha(s) := x_{\lfloor sN\rfloor}^{N,\alpha}$ converges uniformly on such compact intervals to the solution of

$$x'_\alpha(s) = -1 - \frac{x_\alpha(s)}{1-2s}, \qquad x_\alpha(0) = \alpha. \tag{S1.9}$$

Proof. The normalised jumps are bounded by $2/N$. On the positive region the indicator is one, and the displayed conditional expectation gives the drift $-1-x/(1-2s)+o(1)$ after multiplying by N. This drift is locally Lipschitz on $[c,1]\times[0,T]$ for every $c>0$, and the initial condition satisfies $x_0^{N,\alpha} \to \alpha$. Theorem S1.21 applies. $\square$

Corollary S1.23 (Closed form, frontier, and acceleration profile)**.** *The solution of* (S1.9) *is*

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

It hits zero at

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

and the corresponding acceleration profile is

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

In particular,

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

Proof. Solve the linear ODE in (S1.9) and impose $x_{\alpha_\star(s)}(s) = 0$. □

Remark S1.24 (A-side versus mirror E-side). The A-side acceleration chain and the mirror E-side chain are different economic objects, but after state-timing normalisation they share the same fluid drift. The same $\alpha_\star$ and $h_\star$ therefore appear in both the baseline frontier convergence and the $\rho = 0.3$ frontier envelope.

S1.7 Uniform MX static-block identity on $[0, 0.3]$

Theorem 4.2 asserts that the realised ME − MX gap equals the shielding gap $\Delta(\rho) = V_F(\rho) - \mathcal{E}_F(\rho)$ and is strictly increasing on $[0, 0.3]$. Pinning this identification at the endpoints $\rho \in \{0, 0.3\}$ is direct; the single missing ingredient for the interior is a *uniform* MX static-block limit, $N^{-1}\mathrm{PC}_N^{\mathrm{MX}} \to V_G(\rho) = \mu_E - V_F(\rho)$ for every $\rho \in [0, 0.3]$. We supply it by a static counting argument with a fixed safety buffer, avoiding any dynamic control of the MX ratio path.

Reservoir geometry. Work in latent coordinates $X := Z_A$, $Y := Z_E$, with $\mathsf{m}(t) = \ell + (\nu - \ell)\Phi(t)$, $\ell = 0.1$, $\nu = 10$, so $A = \mathsf{m}(X)$, $E = \mathsf{m}(Y)$ are $\mathrm{Unif}[\ell, \nu]$. The static high-A Farmer threshold is

$$q_F = \sqrt{50.005} = 7.071421\ldots, \qquad \mathbb{E}[A\,\mathbf{1}\{A \geq q_F\}] = \tfrac{1}{2}\mu_A = 2.525,$$

the defining property of q_F (the high-A block costs exactly Farmer's parity budget). Fix the buffer $\varepsilon_0 = 0.05$ and set

$$q_{\varepsilon_0} = q_F - \varepsilon_0 = 7.021421\ldots, \qquad \phi_{\varepsilon_0} = \Pr(A \geq q_{\varepsilon_0}) = 0.300867\ldots.$$

Define the buffered safe MX reservoir

$$S_{\rho,\varepsilon_0} = \Big\{A < q_{\varepsilon_0},\ \tfrac{E}{A} > \tfrac{\nu}{q_{\varepsilon_0}}\Big\}.$$

In latent coordinates $A < q_{\varepsilon_0} \iff X < z_{\varepsilon_0}$ with $z_{\varepsilon_0} = \Phi^{-1}\big((q_{\varepsilon_0} - \ell)/(\nu - \ell)\big) = 0.521910\ldots$, and, given $X = x$, the ratio condition reads $Y > \eta_{\varepsilon_0}(x)$ where

$$\eta_{\varepsilon_0}(x) = \Phi^{-1}\left(\frac{(\nu/q_{\varepsilon_0})\,\mathsf{m}(x) - \ell}{\nu - \ell}\right).$$

Hence the buffered reservoir mass is the one-dimensional integral

$$p_{\varepsilon_0}(\rho) := \Pr_\rho(S_{\rho,\varepsilon_0}) = \int_{-\infty}^{z_{\varepsilon_0}} \overline{\Phi}\left(\frac{\eta_{\varepsilon_0}(x) - \rho x}{\sqrt{1-\rho^2}}\right) \varphi(x)\, dx. \tag{S1.10}$$

Lemma S1.25 (Uniform MX ratio-reservoir separation on $[0, 0.3]$)**.** *Fix $\varepsilon_0 = 0.05$, $q_{\varepsilon_0} = q_F - \varepsilon_0$, and S_{ρ,ε_0} as above. A verified interval calculation (Lemma S1.26) gives*

$$\inf_{\rho \in [0,0.3]} \big[\Pr_\rho(S_{\rho,\varepsilon_0}) - \Pr(A \geq q_{\varepsilon_0})\big] > 0.01.$$

Consequently, under MX *against a Rival-Blind max-A Farmer, with parity budgets and full leakage, Green purchases no plot on Farmer's active high-A frontier, with probability tending to one. Hence, up to $o(N)$ terminal filler labels, Farmer's asymptotic capture under* MX *is the* *static descending-A block, and for each* $\rho \in [0, 0.3]$,

$$\lim_{N\to\infty} N^{-1}\mathrm{PC}_N^{\mathrm{MX}} = \mu_E - V_F(\rho) = V_G(\rho).$$

Proof. Throughout, parity budgets are the *realised* budgets

$$B_F^0 = B_G^0 = \tfrac{1}{2}\textstyle\sum_{i=1}^N A_i, \qquad \text{so} \qquad B_F^0/N = B_G^0/N = \tfrac{1}{2}\bar{A}_N = \mu_A/2 + O_p(N^{-1/2}).$$

The comparisons below are population per-plot comparisons with fixed positive margins, so they hold for the realised budgets with probability tending to one.

Piece 1: ratio domination (exact). If $A \geq q_{\varepsilon_0}$ then $E/A \leq \nu/A \leq \nu/q_{\varepsilon_0}$, since $E \leq \nu$; every plot of the band $\{A \geq q_{\varepsilon_0}\}$ has ratio at most ν/q_{ε_0}. Every reservoir plot has ratio strictly above ν/q_{ε_0}, hence strictly above every plot of the band. Since MX selects in strictly decreasing E/A order, it exhausts all available reservoir plots before selecting any plot with $A \geq q_{\varepsilon_0}$. (Ratio ties have probability zero.)

Piece 2: reservoir count margin. By Lemma S1.26 there is $\delta > 0.01$ with $p_{\varepsilon_0}(\rho) \geq \phi_{\varepsilon_0} + \delta$ on $[0, 0.3]$. Since empirical counts deviate from their population values by $O_p(\sqrt{N}) = o_p(N)$, with probability tending to one

$$\big|S_{\rho,\varepsilon_0}^{(N)}\big| > \big|\{i : A_i \geq q_{\varepsilon_0}\}\big|,$$

i.e. the safe reservoir contains strictly more plots than the entire $\{A \geq q_{\varepsilon_0}\}$ band.

Piece 3: Farmer's high-A frontier is blocked above q_{ε_0}. The population A-spend required to clear the buffered band is

$$\mathbb{E}[A\,\mathbf{1}\{A \geq q_{\varepsilon_0}\}] = 2.560588\ldots \;>\; 2.525 = \mu_A/2,$$

so

$$\sum_{i:A_i \geq q_{\varepsilon_0}} A_i - B_F^0 = (2.560588 - 2.525)N + O_p(\sqrt{N}) > 0$$

with probability tending to one. Farmer therefore cannot clear $\{A \geq q_{\varepsilon_0}\}$: its descending high-$A$ frontier is blocked strictly above q_{ε_0}. Once the next frontier plot is unaffordable, Farmer's residual budget is below ν, so any terminal highest-affordable filler purchases have total cost at most ν, cardinality at most ν/ℓ, and ecological value $O(1)$; they are $o(N)$ in normalised Purchased Conservation.

Piece 4: Green affordability. Before Farmer's frontier is blocked Green makes at most $|\{i : A_i \geq q_{\varepsilon_0}\}| + O(1) = \phi_{\varepsilon_0} N + o_p(N)$ purchases, drawn by Pieces 1–2 from the reservoir, hence each of A-cost below q_{ε_0}. Green's pre-frontier spend is therefore at most

$$\phi_{\varepsilon_0}\, q_{\varepsilon_0} N + o_p(N) = 2.112511\ldots N + o_p(N), \qquad \text{while} \qquad B_G^0 = 2.525N + O_p(\sqrt{N}),$$

an order-N affordability buffer. Every remaining reservoir plot is thus individually affordable throughout the pre-frontier phase, so MX keeps selecting from the reservoir rather than from Farmer's high-A frontier.

Piece 5: bounded-cost cleanup. After its frontier is blocked Green continues from the remaining complement. The realised board cost differs from $N\mu_A$ by $O_p(\sqrt{N})$, and prices and values are bounded in $[\ell, \nu]$, so any realised-versus-population budget mismatch changes normalised PC by $o_p(1)$. The terminal filler set of Piece 3 and any final affordability boundary contain $O(1)$ plots, hence contribute $o(N)$. Apart from these $o(N)$ labels Farmer captures the static descending-A block; since $N^{-1}\sum_i E_i \to \mu_E$,

$$N^{-1}\mathrm{PC}_N^{\mathrm{MX}} \to \mu_E - V_F(\rho) = V_G(\rho).$$

□

Lemma S1.26 (Reservoir-mass certificate, $\varepsilon_0 = 0.05$)**.** *With* $\varepsilon_0 = 0.05$*,* $q_{\varepsilon_0} = 7.021421\ldots$*,* $\phi_{\varepsilon_0} = 0.300867\ldots$*, and* $p_{\varepsilon_0}(\rho)$ *of* (S1.10)*, an outward-rounded interval sweep over the* 60 *uniform boxes* $[0, 0.005], \ldots, [0.295, 0.300]$ *certifies*

$$\inf_{\rho \in [0,0.3]} [p_{\varepsilon_0}(\rho) - \phi_{\varepsilon_0}] \;\geq\; 0.0148 \;>\; 0.01,$$

with the binding box $[0.295, 0.300]$.

Certificate. The certified lower bound on p_{ε_0} per box reuses, verbatim, the primitives of the diagonal certificate (Lemma S1.31): IEEE outward rounding, the Abramowitz–Stegun Φ enclosure with its 7.5×10^{-8} pad, the verified inverse-normal brackets, and the nonnegative midpoint quadrature. On each ρ-box and x-cell the integrand of (S1.10) is enclosed by monotone interval bounds: η_{ε_0} at its cell-maximum, ρx and $\sqrt{1-\rho^2}$ at the box-extreme values that maximise the argument, $\overline{\Phi}$ therefore at its minimum, and φ at its cell-minimum; q_{ε_0} is pinned at a rigorous lower bound (shrinking the region and the mass while enlarging ϕ_{ε_0}), and the $x < -8$, $(z_{\varepsilon_0}^-, z_{\varepsilon_0})$, and ratio-singular boundary cells are dropped, all nonnegatively. The script (`reservoir_certificate.py`) and machine report (`reservoir_certificate_report.json`) record $\inf_{[0,0.3]} p_{\varepsilon_0} \geq 0.31572$ at box $[0.295, 0.300]$, whence $\inf(p_{\varepsilon_0} - \phi_{\varepsilon_0}) \geq 0.01485$. □

Consequence for Theorem 4.2. Granting Lemmas S1.25 and S1.26, the MX static-block limit $N^{-1}\mathrm{PC}_N^{\mathrm{MX}} \to V_G(\rho)$ holds uniformly on $[0, 0.3]$. Combined with the ME-side diagonal collapse and the cost-equal exchange (Lemmas S1.27, S1.28, S1.29), the realised gap $N^{-1}\mathbb{E}[\mathrm{PC}_N^{\mathrm{ME}} - \mathrm{PC}_N^{\mathrm{MX}}]$ equals the shielding gap $\Delta(\rho)$ on all of $[0, 0.3]$, not merely at the endpoints, which is the identification asserted by Theorem 4.2:

$$\Delta(\rho) = \lim_{N\to\infty} \frac{1}{N}\mathbb{E}[\mathrm{PC}_N^{\mathrm{ME}} - \mathrm{PC}_N^{\mathrm{MX}}] \text{ exists and is strictly increasing on } [0, 0.3].$$

Scope. The reservoir condition is specific to $[0, 0.3]$: the margin $p_{\varepsilon_0}(\rho) - \phi_{\varepsilon_0}$ is decreasing and the unbuffered mass crosses ϕ near $\rho \approx 0.471$, so the argument does not reach the high-correlation regime and does not by itself deliver a global $\rho \to 1$ theorem. This suffices for the Bolivia-relevant range, in which the estimated $\rho \approx 0.29$ lies.

S1.8 Positive-correlation extension at $\rho = 0.3$: diagonal shielding with frontier-certificate cross-check

This subsection proves the same Budget-World theorem object as Appendix S1.5, but under the Gaussian-copula law

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

where X, ε are independent standard normals. Prices equal A_i, budgets are parity budgets, leakage is full, Farmer is the naïve max-A rule moving first, ME is affordable max-E, and MX is affordable max-R, $R = E/A$. Purchased Conservation is

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

Because E is a strictly increasing transform of the latent Y, the ME rule is equivalently max-Y on the latent scale.

We give two proofs. The primary argument (Section S1.8.1) is a *diagonal-shielding* reduction: in latent coordinates the value-first active frontier collapses onto the diagonal, Farmer's capture becomes a cost-pinned static wedge, and the ME − MX gap is a cost-equal exchange across that diagonal. This yields an *exact* value for the asymptotic gap and a verified-arithmetic monotonicity certificate on $[0, 0.3]$, and replaces the independence symmetry used at $\rho = 0$ without the frontier integration. The original argument, retained below as an independent cross-check (Sections S1.8.3–S1.8.8), proves the same dominance through three packages: (i) a static MX-side safe-frontier identity, (ii) a distribution-free prefix lower bound for ME, and (iii) the global ME frontier/value certificate with a stopping-time affordability bootstrap. The two routes agree numerically; the diagonal-shielding value $\Delta(0.3) = 0.293565\ldots$ strictly improves the frontier certificate's lower bound $0.198653\ldots$

S1.8.1 Diagonal shielding and correlation monotonicity

The positive-correlation case has more structure than the validated frontier certificate of Section S1.8.8 exposes. In latent coordinates $X := Z_A$, $Y := Z_E$, the value-first active frontier collapses onto the diagonal, Farmer's realised capture becomes a static *diagonal wedge*, and the shielding gap $\Delta(\rho) = V_F(\rho) - \mathcal{E}_F(\rho)$ of Lemma S1.5 reduces to a cost-equal exchange across that diagonal. This yields an *exact* value for the $\rho = 0.3$ gap (Theorem S1.30), a one-line structural proof of nonnegativity at every ρ for which the diagonal-wedge representation holds, in particular throughout $[0, 0.3]$ (Lemma S1.29), and a boundary-integral derivative that certifies monotonicity on the empirically relevant range $[0, 0.3]$ (Lemma S1.31).

Throughout, write the common marginal value map

$$\mathsf{m}(t) := \ell + (\nu - \ell)\Phi(t), \qquad \mathsf{m}'(t) = (\nu - \ell)\varphi(t),$$

so that $A = \mathsf{m}(X)$ and $E = \mathsf{m}(Y)$; the diagonal $Y = X$ is exactly the contour $\mathsf{m}(Y) = \mathsf{m}(X)$, i.e. $E = A$. Let $\zeta(t) := \Phi^{-1}(1 - t)$ be the tail-mass quantile (the latent level after removing mass t; distinct from the value-unit frontier q_F), let $\sigma := \sqrt{1-\rho^2}$, let f_ρ be the standard bivariate-normal density with correlation ρ, and set

$$\alpha_\rho := \sqrt{\tfrac{1-\rho}{1+\rho}} = \frac{1-\rho}{\sigma}, \qquad \Phi(\alpha_\rho x) = \Pr(Y \le X \mid X = x).$$

Lemma S1.27 (Diagonal collapse of the value-first frontier)**.** *On the active phase $[0,\tau]$ the value-first fluid frontier system*

$$\dot u = \frac{1}{\Phi((\zeta(v) - \rho\,\zeta(u))/\sigma)}, \qquad \dot v = \frac{1}{\Phi((\zeta(u) - \rho\,\zeta(v))/\sigma)}, \qquad u(0) = v(0) = 0,$$

admits the diagonal solution $u(s) = v(s) = r(s)$ with

$$\text{(S1.11)} \qquad \dot r = \frac{1}{\Phi(\alpha_\rho\,\zeta(r))}, \qquad \textit{equivalently} \qquad \frac{ds}{dr} = \Phi(\alpha_\rho\,\zeta(r)),$$

and any admissible solution coincides with it: $u \equiv v$.

Proof. The system is symmetric under $(u,v) \mapsto (v,u)$. On $u = v = r$ the two right-hand sides agree and equal $1/\Phi(\zeta(r)(1-\rho)/\sigma) = 1/\Phi(\alpha_\rho \zeta(r))$, giving (S1.11); at $\rho = 0$ this integrates to $r(s) = 1 - \sqrt{1-2s}$, recovering the baseline symmetry of Section S1.5. For uniqueness, set $w := (u-v)^2$. Where $u, v > 0$ and, say, $u > v$, the quantile ζ is decreasing, so $\zeta(u) < \zeta(v)$ and

$$(\zeta(v) - \rho\,\zeta(u)) - (\zeta(u) - \rho\,\zeta(v)) = (1+\rho)(\zeta(v) - \zeta(u)) > 0;$$

since Φ is increasing and $1/\Phi$ decreasing, $\dot u < \dot v$, whence $(u-v)(\dot u - \dot v) \le 0$ and $\dot w \le 0$. For admissible solutions—continuous at the entrance with $u(0^+) = v(0^+) = 0$, absolutely continuous on $(0,\tau]$, entering the interior for $s > 0$ (Lemma S1.48)—monotonicity of w on (ε, t) gives $w(t) \le w(\varepsilon)$ for all $0 < \varepsilon < t$; letting $\varepsilon \downarrow 0$ and using continuity gives $w(t) \le w(0^+) = 0$, so $w \equiv 0$. The non-Lipschitz behaviour of ζ' at the corner is therefore immaterial: the squared-gap contraction closes uniqueness once admissibility is granted. $\square$

Lemma S1.28 (Wedge representation of Farmer capture)**.** *Let x_ρ be the unique root of the budget equation*

$$\int_{x_\rho}^{\infty} \mathsf{m}(x)\,\Phi(\alpha_\rho x)\,\varphi(x)\,dx = \tfrac{1}{2}\mu_A. \tag{S1.12}$$

Then Farmer's value-first ecological capture is the static wedge integral

(S1.13)

$$\mathcal{E}_F(\rho) = \mathbb{E}_\rho\big[\mathsf{m}(Y)\,\mathbf{1}\{X \geq x_\rho,\ Y \leq X\}\big] = \int_{x_\rho}^{\infty}\Big[\ell\,\Phi(\alpha_\rho x) + (\nu-\ell)\,\Phi_2\Big(\alpha_\rho x, \tfrac{\rho x}{\sqrt{2-\rho^2}}; -\tfrac{\sqrt{1-\rho^2}}{\sqrt{2-\rho^2}}\Big)\Big]\varphi(x)\,dx.$$

Proof. By Lemma S1.27 both frontiers descend through the common latent level $\zeta(r(s))$ on $[0,\tau]$: Farmer removes the top-X survivors and Green the top-Y survivors at the same threshold. A plot (X, Y) with $X \geq x_\rho := \zeta(r(\tau))$ is reached by Farmer when the descending frontier meets its X-level; at that instant Green's frontier equals X, so the plot is still unsold iff $Y \leq X$. Hence Farmer's captured set is exactly $F_\rho = \{X \geq x_\rho,\ Y \leq X\}$, its A-cost is $\mathbb{E}[\mathsf{m}(X)\mathbf{1}_{F_\rho}] = \frac{1}{2}\mu_A$ (equation (S1.12)), and its E-content is the first expression in (S1.13). For the closed form, condition on $X = x$: then $Y \sim N(\rho x, \sigma^2)$ and $\Pr(Y \leq x \mid X = x) = \Phi(\alpha_\rho x)$, giving the ℓ-term. For the slope term, introduce $W \sim N(0,1)$ independent of Y; since $\Phi(Y) = \Pr(W \leq Y \mid Y)$,

$$\mathbb{E}\big[\Phi(Y)\mathbf{1}\{Y \leq x\} \mid X = x\big] = \Pr\left(W \leq Y,\ Y \leq x \mid X = x\right).$$

Writing $Y = \rho x + \sigma Z$ with $Z \sim N(0,1)$: $\{Y \leq x\} = \{Z \leq \alpha_\rho x\}$, and $\{W \leq Y\} = \{(W - \sigma Z)/\sqrt{2-\rho^2} \leq \rho x/\sqrt{2-\rho^2}\}$. The pair $(Z, (W-\sigma Z)/\sqrt{2-\rho^2})$ is standard bivariate normal with correlation $-\sigma/\sqrt{2-\rho^2}$, yielding the Φ_2 term and hence (S1.13). This is the exact replacement for the frontier/Picard object of Section S1.8.8: one scalar threshold and one one-dimensional integral.

It remains to confirm that Green is still active at Farmer exhaustion, so that the common-frontier picture is valid up to τ. Green's active captures form the reflected wedge $G_\rho = \{Y \geq x_\rho,\ X < Y\}$, on which it pays the price $\mathsf{m}(X)$. By exchangeability of (X, Y),

$$\mathbb{E}\big[\mathsf{m}(X)\mathbf{1}\{Y \geq x_\rho,\ X < Y\}\big] = \mathbb{E}\big[\mathsf{m}(Y)\mathbf{1}\{X \geq x_\rho,\ Y < X\}\big] = \mathcal{E}_F(\rho),$$

and on F_ρ one has $\mathsf{m}(Y) < \mathsf{m}(X)$ off the diagonal, so this is strictly below $\mathbb{E}[\mathsf{m}(X)\mathbf{1}_{F_\rho}] = \frac{1}{2}\mu_A$. Hence Green's active A-spend through Farmer exhaustion is $\mathcal{E}_F(\rho) < \frac{1}{2}\mu_A$: Green never binds before τ, and the affordability bootstrap is analytic, not numerical. $\square$

Lemma S1.29 (Diagonal-shielding exchange and nonnegativity)**.** *With z_F the static high-A cutoff ($\int_{z_F}^{\infty} \mathsf{m}(x)\varphi(x)\,dx = \frac{1}{2}\mu_A$) and*

$$H_\rho := \{X \geq z_F,\ Y > X\}, \qquad L_\rho := \{x_\rho \leq X < z_F,\ Y \leq X\},$$

the static high-A block $\{X \geq z_F\}$ and the Farmer wedge F_ρ spend the same A-budget, so $\mathbb{E}[\mathsf{m}(X)\mathbf{1}_{H_\rho}] = \mathbb{E}[\mathsf{m}(X)\mathbf{1}_{L_\rho}]$, and

$$\Delta(\rho) = V_F(\rho) - \mathcal{E}_F(\rho) = \mathbb{E}\big[(\mathsf{m}(Y) - \mathsf{m}(X))\mathbf{1}_{H_\rho}\big] + \mathbb{E}\big[(\mathsf{m}(X) - \mathsf{m}(Y))\mathbf{1}_{L_\rho}\big] \geq 0. \tag{S1.14}$$

For every ρ at which the diagonal-wedge representation (Lemma S1.28) has been established the bound is nonnegative; in particular it holds on $[0, 0.3]$, the range covered by the entrance lemma S1.48. If the entrance/admissibility argument of Lemma S1.27 is extended pointwise to every fixed $\rho < 1$, the bound holds on all of $[0, 1)$. Moreover $\Delta(\rho) \to 0$ as $\rho \uparrow 1$: mass concentrates on the diagonal $Y = X$, where both exchange surpluses vanish.

Proof. Both $\{X \geq z_F\}$ and $F_\rho = \{X \geq x_\rho, Y \leq X\}$ have A-cost $\frac{1}{2}\mu_A$; subtracting their common overlap $\{X \geq z_F, Y \leq X\}$ leaves H_ρ on the static side and L_ρ on the wedge side with equal A-

budget, i.e. $\mathbb{E}[\mathsf{m}(X)\mathbf{1}_{H_\rho}] = \mathbb{E}[\mathsf{m}(X)\mathbf{1}_{L_\rho}]$. Since $V_F = \mathbb{E}[\mathsf{m}(Y)\mathbf{1}\{X \geq z_F\}]$ and $\mathcal{E}_F = \mathbb{E}[\mathsf{m}(Y)\mathbf{1}_{F_\rho}]$ (Lemma S1.28), $\Delta = \mathbb{E}[\mathsf{m}(Y)\mathbf{1}_{H_\rho}] - \mathbb{E}[\mathsf{m}(Y)\mathbf{1}_{L_\rho}]$; adding and subtracting the equal A-budgets gives (S1.14). On H_ρ, $Y > X$ so $\mathsf{m}(Y) > \mathsf{m}(X)$; on L_ρ, $Y \leq X$ so $\mathsf{m}(X) \geq \mathsf{m}(Y)$, since m is the same increasing map on both coordinates. Both terms are nonnegative. □

Theorem S1.30 (Exact value-first dominance at $\rho = 0.3$). *For the $\rho = 0.3$ Budget-World object,*

$$\lim_{N\to\infty} \frac{1}{N}\,\mathbb{E}[\mathrm{PC}_N^{\mathrm{ME}} - \mathrm{PC}_N^{\mathrm{MX}}] = \Delta(0.3) = V_F(0.3) - \mathcal{E}_F(0.3) = 1.784882\ldots - 1.491317\ldots = 0.293565\ldots > 0.$$

Proof. The limit is identified in four steps. (i) The MX side has the closed-form static-block limit $V_G(0.3)$ (Proposition S1.58), so by the shielding identity (Lemma S1.5) every subsequential gap limit equals $V_F(0.3) - \mathcal{E}_F^\star(0.3)$. (ii) Lemma S1.27 pins a unique deterministic value-first fluid path, so $N^{-1}\mathrm{PC}_N^{\mathrm{ME}}$ has a single limit, not merely a subsequential one; the subsequential hedge $\mathcal{E}_F^\star$ of Lemma S1.5 collapses to $\mathcal{E}_F(0.3)$. (iii) Lemma S1.28 evaluates that unique Farmer capture: with $\alpha_{0.3} = \sqrt{7/13}$, the budget root of (S1.12) is $x_{0.3} = 0.181411\ldots$, and (S1.13) gives $\mathcal{E}_F(0.3) = 1.491317\ldots$. (iv) Boundedness of $E \in [\ell, \nu] = [0.1, 10]$ upgrades the in-probability / subsequential convergence to convergence of normalised expectations. With $V_F(0.3) = 1.784882\ldots$ from Lemma S1.3, the gap is $\Delta(0.3) = 0.293565\ldots$. Nonnegativity is structural (Lemma S1.29); the displayed value is the exact strengthening of the validated lower bound $0.198653\ldots$ of Corollary S1.61. □

Lemma S1.31 (Boundary derivative and monotonicity on $[0, 0.3]$). *The shielding gap is differentiable in ρ with*

$$\Delta'(\rho) = Z_\rho - C_\rho - R_\rho - x'_\rho\, M_\rho, \tag{S1.15}$$

where, with f_ρ the correlation-ρ standard bivariate-normal density,

$$Z_\rho = \int_{-\infty}^{\infty} \mathsf{m}'(y) f_\rho(z_F, y)\, dy = V_F'(\rho), \quad C_\rho = \int_{-\infty}^{x_\rho} \mathsf{m}'(y) f_\rho(x_\rho, y)\, dy, \quad R_\rho = \int_{x_\rho}^{\infty} \mathsf{m}'(x) f_\rho(x, x)\, dx,$$

$$M_\rho = \int_{-\infty}^{x_\rho} (\mathsf{m}(x_\rho) - \mathsf{m}(y)) f_\rho(x_\rho, y)\, dy.$$

Since $x_\rho' < 0$ on $[0, 0.3]$, the term $-x_\rho' M_\rho$ is positive, and a verified interval-arithmetic sweep (below) gives

$$\Delta'(\rho) > 0.20 \qquad \textit{for all } \rho \in [0, 0.3];$$

in particular Δ is strictly increasing on $[0, 0.3]$ and $\Delta(0.3) > \Delta(0) = 0.216135\ldots$

Proof. Differentiate the exchange form (S1.14). Inside each region the integrand is additively separable in (X, Y) ($\mathsf{m}(Y) - \mathsf{m}(X)$ on H_ρ, $\mathsf{m}(X) - \mathsf{m}(Y)$ on L_ρ), so by Price's theorem $\partial_\rho f_\rho = \partial_x \partial_y f_\rho$ the interior contribution vanishes and the ρ-derivative is carried entirely by the region boundaries: the fixed edge $X = z_F$ gives Z_ρ (equal to $V_F'(\rho)$ after one integration by parts in y); the fixed edge $X = x_\rho$ gives $-C_\rho$; the diagonal edge $Y = X$ gives $-R_\rho$; and the motion of the cutoff contributes $-x_\rho' M_\rho$, with $x_\rho' = -\partial_\rho B / \partial_x B$ from (S1.12), $B(\rho, x) = \int_x^\infty \mathsf{m}(t) \Phi(\alpha_\rho t) \varphi(t)\, dt$, and $\partial_x B = -\mathsf{m}(x) \Phi(\alpha_\rho x) \varphi(x) < 0$ and $\partial_\rho B(\rho, x) = \int_x^\infty \mathsf{m}(t)\, \varphi(\alpha_\rho t)\, \alpha_\rho'\, t\, \varphi(t)\, dt < 0$ (since $\alpha_\rho' = -1/[\alpha_\rho (1+\rho)^2] < 0$ and the integrand is otherwise positive on $t \geq x_\rho > 0$), so $x_\rho' = -\partial_\rho B / \partial_x B < 0$ on $[0, 0.3]$ (and $x_\rho > 0$ there, since the certified precheck gives $B(0.3, 0) > \frac{1}{2}\mu_A$). A verified interval-arithmetic sweep over the 60 subintervals $[0, 0.005], [0.005, 0.010], \ldots, [0.295, 0.300]$ encloses the four boundary terms over each box (using monotone endpoint bracketing $x_\rho \in [x_{\rho_U}, x_{\rho_L}]$ and the magnitude form $|x_\rho'| = (-\partial_\rho B)/(-\partial_x B)$) and evaluates the one-sided lower bound

$$\Delta'(\rho) \;\geq\; Z_{\min} - C_{\max} - R_{\max} + |x_\rho'|_{\min} M_{\min}$$

on each box. The worst certified box is $[0.295, 0.300]$, with lower bound $0.222487\ldots$; hence $\Delta'(\rho) > 0.20$ for all $\rho \in [0, 0.3]$, and Δ is strictly increasing there, so $\Delta(0.3) > \Delta(0)$. For orientation, finer point evaluations give $\Delta'(0) = 0.274349\ldots$ and $\Delta'(0.3) = 0.234605\ldots$, the smallest observed point value being at the right endpoint, though we do not establish that Δ' is monotone; the certified statement uses only the swept lower bound. The certificate uses the same outward-rounded IEEE arithmetic, Abramowitz–Stegun normal-CDF enclosure with 7.5×10^{-8} padding, verified inverse-normal brackets, and midpoint-plus-remainder quadrature as the $\rho = 0.3$ fluid-limit certificate of Section S1.8.8. □

Remark S1.32 (Scope: shielding gap versus the global ME − MX gap). Lemma S1.29 signs the shielding gap $\Delta(\rho) = V_F - \mathcal{E}_F$ at every ρ. Identifying this with the realised ME − MX Budget-World gap uses the MX static-block limit, which the ratio-reservoir argument of Appendix S1.7 establishes uniformly on $[0, 0.3]$; combined with the monotonicity of the shielding object certified in the present subsection, the realised gap is strictly increasing on $[0, 0.3]$. The simulated purchased-conservation gap reported in the main text is hump-shaped, peaking near $\rho \approx 0.70$ (where the diagonal term R_ρ, the only term that grows without bound as $\rho \to 1$, overtakes $Z_\rho - x'_\rho M_\rho$); a global high-correlation statement would additionally require extending the MX safe-frontier analysis into the regime where the top-ratio block begins to overlap the high-A block, and is left as a separate extension.

S1.8.2 Frontier/fluid-limit cross-check

The remainder of this subsection develops the original frontier/value certificate. It is logically self-contained and independent of the diagonal-shielding argument above; it is retained because it certifies the same $\rho = 0.3$ dominance by a different route (an explicit fluid-limit ODE with a validated Picard step-tube) and because the diagonal-shielding theorem cites its MX-side threshold identity (Proposition S1.58) and entrance-uniqueness lemma (Lemma S1.48).

S1.8.3 Proof target

Set $m = 0.10$. Proposition S1.58 gives the MX threshold limit

$$N^{-1}\mathrm{PC}_N^{\mathrm{MX}} \to \mathbb{E}[E\,\mathbf{1}\{A \le q_F\}],$$

while the prefix bound gives

$$\liminf_N N^{-1} \sum_{t \le mN} E(g_t^{\mathrm{ME}}) \ge ME(m) = 0.901.$$

It therefore suffices to prove

$$\liminf_N N^{-1}\mathbb{E}\left[\sum_{t=mN+1}^{\lfloor \kappa_{\mathrm{low}} N \rfloor} E(g_t^{\mathrm{ME}})\right] > \mathbb{E}[E\,\mathbf{1}\{A \le q_F\}] - 0.901.$$

The certified upper bound on the right side is 2.36413. The global fluid-limit and affordability bootstrap in Sections S1.8.7–S1.8.8 prove this inequality.

S1.8.4 Static MX side and the prefix wedge

Let $r_s(0.3)$ be the ratio cutoff satisfying

$$\Pr(R \ge r_s(0.3)) = s.$$

The certified ratio-cut intervals are

$$\text{(S1.16)} \qquad r_{0.10}(0.3) \in [3.62406450, 3.62406459], \qquad r_\phi(0.3) \in [1.50517830, 1.50517845].$$

All inequalities below use the conservative endpoints of these intervals.

Lemma S1.33 (MX-side Green affordability, $\rho = 0.3$)**.** *Let*

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

Then $C_R^{(0.3)}$ is continuous and strictly decreasing, and the archived outward-rounded interval quadrature gives

$$C_R^{(0.3)}(r_\phi(0.3)) \in [0.679, 0.685] < \frac{\mu_A}{2} = 2.525. \tag{S1.17}$$

Consequently, for some $\delta_{\text{aff}} > 0$ and $\varepsilon_0 \geq 1.83$,

$$C_R^{(0.3)}(r_{\phi+\delta}(0.3)) \leq \frac{\mu_A}{2} - \varepsilon_0, \qquad 0 \leq \delta \leq \delta_{\text{aff}}.$$

Proof. Absolute continuity of the Gaussian-copula density on the compact support gives continuity and strict monotonicity. The interval (S1.17) is the certified latent-coordinate quadrature of the ratio half-space $\{R \geq r\}$, evaluated at the endpoints of (S1.16) with the same one-sided Abramowitz–Stegun Φ-bounds, inverse-normal brackets, and outward-rounded arithmetic described in Appendix S1.8.8; it is the certificate key `mx_affordability_rho03`. The uniform margin then gives δ_{aff} by continuity of r_s and $C_R^{(0.3)}$. ☐

Lemma S1.34 ($\rho = 0.3$ safe-frontier buffer)**.** *The strict safe-frontier inequality*

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

holds. Hence there exist $\delta_{\text{sf}}, \eta_{\text{sf}} > 0$ 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 $\Omega_N^{\mathrm{sf}}(\delta_{\mathrm{sf}}, \eta_{\mathrm{sf}})$ has probability tending to one. On this event, the initial top-$(\phi + \delta_{\mathrm{sf}})N$ ratio block is disjoint from the first $\lfloor (\phi + \delta_{\mathrm{sf}})N \rfloor$ items of the descending A-order.

Proof. Using the lower endpoint in (S1.16),

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

Continuity of $s \mapsto r_s(0.3)$ and the uniform A-quantile map gives the buffer constants. Proposition S1.14 gives the empirical cutoff event with probability tending to one, and the corner-disjointness argument is the same as in Lemma S1.18. □

Corollary S1.35 (Green is unconstrained through the MX safe-frontier window)**.** *Let $\delta_0 := \min(\delta_{\mathrm{aff}}, \delta_{\mathrm{sf}})$. On an event $\Omega_N^{\mathrm{afford}}$ with probability tending to one,*

$$B_{G,N}^{\mathrm{res}}(t) \geq \frac{\varepsilon_0 N}{2} > \nu \qquad \text{for all } t \leq \lfloor (\phi + \delta_0)N \rfloor$$

and all sufficiently large N. Thus, through this window, MX *Green plays exact* $\arg\max R$ *over the whole survivor set, not merely over the affordable subset.*

Proof. Identical to Corollary S1.19, replacing the baseline ratio-cost function, safe-frontier event, and affordability lemma by Lemmas S1.33 and S1.34. □

Proposition S1.36 (Safe-frontier exactness through ϕ)**.** *Let $q_F = Q_A(1-\phi)$. With probability tending to one:*

(i) MX*'s first $\lfloor mN \rfloor$ Green purchases are the initial static top-ratio portfolio, for $m = 0.10$;*

(ii) through Farmer exhaustion, MX*'s Green purchases agree with the initial descending ratio order up to $o_p(N)$ endpoint labels, and $K_N^F(\mathrm{MX})/N \to \phi$;*

(iii) after Farmer exhaustion, the empirical residual pool converges, for every bounded threshold-continuity test function, to

$$D_{\phi,0.3} := \{(E, A) : A \le q_F,\ E/A < r_\phi(0.3)\}.$$

Hence normalised 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 is $10/r_s(0.3) < Q_A(1-s)$. At $s = 0.10$,

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

and at $s = \phi$,

$$q_F - \frac{10}{\underline{r}_\phi} > 0.42769.$$

By continuity, choose $\delta, \eta > 0$ such that

$$\frac{10}{r_{\phi+\delta}(0.3) - 2\eta} < Q_A(1 - \phi - \delta) - 2\eta. \tag{S1.18}$$

On the event where the empirical ratio and A-cutoffs are within η of their population values and the no-Green descending-A budget crossing time σ_N satisfies $\sigma_N \le (\phi+\delta)N$, the initial top-$(\phi+\delta)N$ ratio block is disjoint from the first $(\phi+\delta)N$ descending-A items. This event has probability tending to one by Proposition S1.14 and the ordinary budget LLN, because $N^{-1}\sigma_N \to \phi$.

On the intersection with the affordability event of Corollary S1.35, induction over $t \le \sigma_N$ gives the exact static path: Farmer buys the t-th item of the initial descending A-order, while MX Green buys the t-th item of the initial descending ratio order. The two blocks are disjoint and Green is unconstrained. At σ_N, Farmer's remaining budget is below 10; since every further purchase costs at least $\ell = 0.1$,

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

Therefore $K_N^F(\mathrm{MX})/N \to \phi$, and Lemma S1.16 replaces any fixed $T < \phi$ by ϕ in normalised sums. Finally, $\{R \ge r_\phi(0.3)\} \subset \{A < q_F\}$; hence the Farmer-active Green block is the static top-ratio block inside $\{A \le q_F\}$, Farmer's active block is the empirical top-A block, and the residual empirical pool is $D_{\phi,0.3}$ up to $o_p(N)$ boundary labels. Thresholded LLNs complete the residual-pool convergence. □

Distribution-free prefix lower bound for ME**.** The prefix comparison uses only the marginal E-law and the fact that, in the first mN rounds with $m = 0.10$, Green can afford the whole survivor set with probability tending to one.

Lemma S1.37 (Alternating order-statistic bound on the ME prefix)**.** *Let $E_{(1)} \ge \cdots \ge E_{(N)}$ be the ecological order statistics. On the event that* ME *Green is unconstrained through its first $\lfloor mN \rfloor$ turns,*

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

Proof. Before Green's t-th move, at most $2t-1$ plots have been removed. Therefore at least one of the top $2t$ ecological plots survives; since Green is unconstrained and chooses max-E, its t-th purchase has value at least $E_{(2t)}$. Sum over t. □

Lemma S1.38 (Adjacent order-statistic spacing under iid uniform marginals)**.** *If $E_1, \ldots, E_N$ are iid $U[\ell, \nu]$, 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

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

Proof. Write $E_i = \ell + dU_i$. Each uniform spacing has mean $1/(N+1)$; hence the nonnegative normalised sum of adjacent gaps has expectation $O(N^{-1})$ and converges to zero by Markov. The identity

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

and the quantile LLN of Proposition S1.14 give the displayed limit. □

Proposition S1.39 (Prefix wedge at $\rho = 0.3$)**.** *Let*

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

and

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

Then

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

and

$$\liminf_{N\to\infty} \frac{1}{N}\mathbb{E}[G^{\mathrm{PC}}_{N,0.10}] \geq 0.901 - 0.59221775 > 0.3087822. \tag{S1.19}$$

Proof. Early affordability applies under the copula law because it uses only $A \leq \nu$ and the normalised-budget LLN: $0.10\nu = 1 < \mu_A/2 = 2.525$. Hence, with probability tending to one, ME is unconstrained through its first $0.10N$ turns. Lemmas S1.37–S1.38 yield

$$N^{-1}\sum_{t\leq 0.10N} E(g_t^{\mathrm{ME}}) \geq \underline{M}_E(0.10) + o_p(1),$$

and direct integration gives

$$\underline{M}_E(0.10) = \tfrac{1}{2}\int_{0.8}^{1}(0.1 + 9.9u)\,du = 0.901.$$

For MX, Proposition S1.36 gives the static top-0.10 ratio block up to $o_p(N)$ boundary labels. The thresholded LLN for the ratio half-space class gives convergence to $M_X(0.10, 0.3)$, and certified quadrature gives the displayed interval (certificate key `mx_prefix_value_rho03`). Boundedness upgrades the in-probability lower bound to the expectation lower bound (S1.19). □

S1.8.5 The ME Farmer-active frontier and value path

Under positive correlation, Farmer deletions are not E-neutral. The proof therefore works on the latent Gaussian scale. Conditional on the surviving X-profile and past Green picks, current survivor Y-scores are independent upper-truncated normals with means ordered by X. The next three statements turn that observation into the Farmer-frontier lower envelope used by the ME value-path certificate.

Lemma S1.40 (Monotone argmax for upper-truncated normals)**.** *Fix a common upper truncation $c \in \mathbb{R} \cup \{+\infty\}$, variance $\sigma^2 > 0$, and means $\mu_1 \le \cdots \le \mu_M$. Let $Y_i^{(c)}$ be independent $N(\mu_i, \sigma^2)$ variables truncated to $(-\infty, c]$. Then*

$$p_i := \Pr(Y_i^{(c)} = \max_j Y_j^{(c)})$$

is nondecreasing in i. Hence any subset containing the k largest means contains the argmax with probability at least k/M.

Proof. Let f^c_μ, F^c_μ be the truncated density and cdf. For $\mu' > \mu$,

$$\frac{d}{dy} \log \frac{f^c_{\mu'}(y)}{f^c_\mu(y)} = \frac{\mu' - \mu}{\sigma^2} \geq 0,$$

so the common-truncation family has monotone likelihood ratio. For $\mu_i \leq \mu_j$, write $r = f_j/f_i$. Since r is nondecreasing,

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

and therefore

$$p_j - p_i = \int_{-\infty}^{c} \{f_j(y)F_i(y) - f_i(y)F_j(y)\} \prod_{\ell \neq i,j} F_\ell(y)\, dy \geq 0.$$

Thus $p_1 \leq \cdots \leq p_M$. The k largest probabilities have average at least the overall average $1/M$, giving the final claim. □

Lemma S1.41 (Conditional law of current survivors)**.** *Fix $T < \phi$, and let Ω^N_T be 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 survivor set immediately before Green's pick, $n_t = |S_t| = N - 2t + 1$, and*

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

Then S_t, n_t, and $(X_i)_{i \in S_t}$ are $\mathcal{G}_t$-measurable. With $c_1 = +\infty$ and $c_t = Y_{g_{t-1}}$ for $t \geq 2$,

$$Y_i \mid \mathcal{G}_t \sim N(\rho X_i, 1 - \rho^2) \ \ \textit{truncated to } (-\infty, c_t], \qquad i \in S_t,$$

independently across $i \in S_t$.

Proof. Conditional on $\mathbf{X}$, the Y_i's are independent with $Y_i | X_i \sim N(\rho X_i, 1-\rho^2)$. Farmer deletions are measurable in $\mathbf{X}$ and past Green indices. On Ω^N_T, Green removes the current $\arg\max Y$, so the previously chosen Green values are strictly decreasing and survival through the first $t-1$ Green turns is equivalent to the single constraint $Y_i \le c_t$. Conditioning on $\mathcal{G}_t$ thus leaves the product density

$$\prod_{i \in S_t} \varphi_{\rho X_i,\, 1-\rho^2}(y_i)\mathbf{1}\{y_i \le c_t\},$$

renormalised coordinatewise. This is the stated independent upper-truncated-normal law. $\square$

Proposition S1.42 (Conditional top-block hit bound)**.** *Fix $T < \phi$ and work on Ω_T^N. For Green turn $t \le \lfloor TN \rfloor$, let $B_t(k) \subseteq S_t$ be the k highest surviving X-values. Then*

$$\Pr(g_t \in B_t(k) \mid \mathcal{G}_t) \ge \frac{k}{n_t} \qquad a.s. \tag{S1.20}$$

If B_α^N is the initial top-$\lfloor \alpha N \rfloor$ A-block and $X_t^{N,\alpha}$ is the number of its survivors in S_t, then, when $X_t^{N,\alpha} > 0$,

$$\Pr(g_t \in B_\alpha^N \mid \mathcal{G}_t) \ge \frac{X_t^{N,\alpha}}{n_t} \qquad a.s. \tag{S1.21}$$

Proof. By Lemma S1.41, conditional survivor scores are independent upper-truncated normals with common variance and means ρX_i, ordered by X_i because $\rho > 0$. Lemma S1.40 gives (S1.20). The survivors of the fixed initial top-αN block are precisely the $X_t^{N,\alpha}$ highest surviving X-values, giving (S1.21). $\square$

For the next proposition, order the initial labels by descending agricultural value A, equivalently by descending latent X, and let S_t be the survivor set immediately before Green's t-th move. Define the empirical Farmer frontier

$$F_t^{N,\mathrm{ME}} := \max\Big\{k \in \{0,\dots,N\} : \{1,\dots,k\}_{A\text{-rank}} \cap S_t = \varnothing\Big\}.$$

Thus $F_t^{N,\mathrm{ME}}/N$ is the largest initial top-A mass that has been completely removed by time t.

The event that the initial top $\lfloor \alpha N \rfloor$ A-block is extinct is equivalent, up to a one-label rounding error, to $F_t^{N,\mathrm{ME}}/N \geq \alpha$. The process $t \mapsto F_t^{N,\mathrm{ME}}$ is nondecreasing.

Proposition S1.43 (Correlated Farmer-frontier lower envelope)**.** *Let*

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

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

$$\Pr\left(\Omega_T^N \cap \left\{ \inf_{0 \leq s \leq T} \frac{F_{\lfloor sN \rfloor}^{N,\mathrm{ME}}}{N} < \alpha_\star(s) - \varepsilon \right\}\right) \to 0. \tag{S1.22}$$

Consequently, if $\Pr(\Omega_T^N) \to 1$, the same lower envelope holds unconditionally with probability tending to one.

Proof. Fix $\alpha \in (0,1)$, and let $X_t^{N,\alpha}$ be the number of survivors from the initial top-$\lfloor \alpha N \rfloor$ A-block before Green turn t. On Ω_T^N, Proposition S1.42 gives

$$p_t^\alpha := \Pr(g_t \in B_\alpha^N \mid \mathcal{G}_t) \geq X_t^{N,\alpha}/n_t.$$

Realise the Green hit indicator as $\mathbf{1}\{U_t \leq p_t^\alpha\}$ using iid uniforms U_t, and define the comparison chain

$$\widetilde{X}_{t+1}^{N,\alpha} = \left(\widetilde{X}_t^{N,\alpha} - 1 - \mathbf{1}\{U_t \leq \widetilde{X}_t^{N,\alpha}/n_t\}\right)^+, \qquad \widetilde{X}_1^{N,\alpha} = X_1^{N,\alpha}.$$

Since $k \mapsto (k - 1 - \mathbf{1}\{u \leq k/n_t\})^+$ is nondecreasing and $p_t^\alpha \geq X_t^{N,\alpha}/n_t$, induction gives the pathwise domination

$$X_t^{N,\alpha} \leq \widetilde{X}_t^{N,\alpha}, \qquad t \leq \lfloor TN \rfloor, \quad \text{on } \Omega_T^N. \tag{S1.23}$$

The scaled comparison chain is exactly the baseline fixed-block depletion chain: jumps are

$O(1/N)$, and its drift on the positive region is

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

with solution

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

Hence Theorem S1.21, Proposition S1.22, and Corollary S1.23 apply to $\widetilde{X}^{N,\alpha}$.

Now choose a grid $0 = u_0 < \cdots < u_m = T$ so fine that $\alpha_\star(u_j) - \alpha_\star(u_{j-1}) < \varepsilon/2$. Let

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

Because $\alpha_j < \alpha_\star(u_j)$, the deterministic path x_{α_j} hits zero before u_j. The comparison-chain LLN and (S1.23) imply that, on Ω_T^N, the actual top-$\alpha_j N$ block is extinct by time $u_j N$ with probability tending to one. Extinction of that block is exactly $F^{N,\mathrm{ME}}_{\lfloor u_j N \rfloor}/N \geq \alpha_j$. A union bound over the finite grid and monotonicity of $F_t^{N,\mathrm{ME}}$ yield (S1.22). □

S1.8.6 Exact dynamic survivor law

Lemma S1.41 conditions on the full latent X-profile and is enough for the frontier envelope. The fluid-limit reduction needs the stronger Markov statement below: conditional on the current thresholds, not on the full profile, survivors are iid from the doubly truncated host.

Lemma S1.44 (Exact dynamic survivor law)**.** *Fix $\rho \in [0,1)$ and work on an event where Farmer and Green remain affordable, hence play exact* $\arg\max X$ *and* $\arg\max Y$*. Let i_t^F, i_t^G be their t-th purchase indices, $x_t = X_{i_t^F}$, and $c_t = Y_{i_t^G}$, with $c_0 = +\infty$. Let $\mathcal{H}_t$ be the full history immediately before Green's t-th move, including the indices and (X,Y)-values of all previous Farmer purchases through t and previous Green purchases through $t-1$. Conditional on $\mathcal{H}_t$,* *the $N-2t+1$ surviving labels are iid from*

$$\nu_{x_t,c_{t-1}} := \mathcal{L}\big((X,Y) \mid X \le x_t,\ Y \le c_{t-1}\big).$$

Proof. Use the standard order-statistics factorisation: if iid pairs (Z_j, W_j) are drawn from a continuous law P and $k = \arg\max Z_j$, then conditional on k and $(Z_k, W_k) = (z, w)$, the remaining pairs are iid from P restricted to $\{Z < z\}$, independently of w. This follows because the conditional density factors as

$$\prod_{j \ne k} \frac{p(z_j, w_j)\mathbf{1}\{z_j < z\}}{P(Z < z)}.$$

Apply this factorisation inductively. At $t = 1$, Farmer's max-X pick leaves the remaining labels iid from the original copula law restricted to $\{X < x_1\}$, i.e. ν_{x_1,c_0}. If before Green turn t survivors are iid from $\nu_{x_t,c_{t-1}}$, Green's max-Y pick leaves survivors iid from ν_{x_t,c_t}; then Farmer's max-X pick leaves survivors iid from ν_{x_{t+1},c_t}. Continuity makes the strict/weak boundary distinction immaterial and rules out ties almost surely. $\square$

The same statement holds after Farmer exhaustion with x_t replaced by the frozen Farmer frontier x_τ: conditional on the post-Farmer history immediately before a Green move, the surviving labels are iid from

$$\mathcal{L}((X,Y) \mid X \le x_\tau,\ Y \le c_{t-1}).$$

Remark S1.45 (The Markov reduction is strong, not marginal)*.* Lemma S1.44 conditions on the full history $\mathcal{H}_t$, but the survivor law depends on that history only through (x_t, c_{t-1}). Thus the normalised state $(x_t, c_{t-1}, b_F(t), b_G(t))$ is Markov, which is the input required for the density-dependent LLN and one-step drift in the next subsection.

S1.8.7 Global fluid limit, validation, and theorem closure

Fix $\rho = 0.3$, set $\sigma = \sqrt{1-\rho^2}$, and write

$$M(x, c; \rho) = \Phi_2(x, c; \rho), \qquad a(x) = 0.1 + 9.9\Phi(x), \qquad e(c) = 0.1 + 9.9\Phi(c).$$

For the agricultural cost of a max-Y Green pick from the doubly truncated host $\{X \leq x, Y \leq c\}$, define

$$g_A(x, c) := \mathbb{E}[A \mid Y = c,\ X \leq x]. \tag{S1.24}$$

Since $X \mid Y = c \sim N(\rho c, \sigma^2)$,

$$g_A(x, c) = 0.1 + 9.9 \frac{\displaystyle\int_{-\infty}^{x} \Phi(u)\sigma^{-1}\varphi\left(\frac{u - \rho c}{\sigma}\right) du}{\Phi\left(\dfrac{x - \rho c}{\sigma}\right)}. \tag{S1.25}$$

Equivalently,

$$g_A(x, c) = 0.1 + 9.9\, \frac{\Phi_2\left(\frac{x-\rho c}{\sigma}, \frac{\rho c}{\sqrt{2-\rho^2}}; -\frac{\sigma}{\sqrt{2-\rho^2}}\right)}{\Phi\left(\frac{x-\rho c}{\sigma}\right)}. \tag{S1.26}$$

During the Farmer-active phase the limiting state $Z(s) = (x(s), c(s), b_F(s), b_G(s))$ solves

$$\dot{x}(s) = -\frac{1}{\varphi(x(s))\,\Phi((c(s) - \rho x(s))/\sigma)}, \tag{S1.27}$$

$$\dot{c}(s) = -\frac{1}{\varphi(c(s))\,\Phi((x(s) - \rho c(s))/\sigma)}, \tag{S1.28}$$

$$\dot{b}_F(s) = -a(x(s)), \qquad \dot{b}_G(s) = -g_A(x(s), c(s)). \tag{S1.29}$$

Let $\tau = \inf\{s : b_F(s) = 0\}$. After Farmer exhaustion, x freezes at $x_\tau = x(\tau)$, and for $s > \tau$,

$$\dot{x}(s) = 0, \qquad \dot{c}(s) = -\frac{1}{\varphi(c(s))\,\Phi((x_\tau - \rho c(s))/\sigma)}, \qquad \dot{b}_G(s) = -g_A(x_\tau, c(s)). \tag{S1.30}$$

The per-turn Green ecological value along the fluid path is $e(c(s))$.

Lemma S1.46 (Survivor mass and one-step drift)**.** *On any stopped horizon on which exact max-X/max-Y play holds, if n_t is the survivor count before Green's t-th move and (x_t, c_{t-1}) are the corresponding frontiers, then for every fixed $T < 1/2$,*

$$\sup_{t \le TN} \left| \frac{n_t}{N} - \Phi_2(x_t, c_{t-1}; \rho) \right| \xrightarrow{p} 0.$$

More precisely the error is $O_p(N^{-1/2}) + O(N^{-1})$ uniformly along the stopped path.

Proof. By Lemma S1.44, after Farmer's t-th purchase and before Green's t-th purchase, the survivor set is the initial lower orthant $\{X \le x_t, Y \le c_{t-1}\}$, up to the boundary label. Thus the difference between n_t/N and the empirical lower-orthant mass is at most $2/N$. Uniform Glivenko–Cantelli convergence for the VC class of bivariate lower orthants gives the result at the random thresholds. □

Lemma S1.47 (One-step drift in latent coordinates)**.** *On compact interior sets with positive remaining budgets, the one-step conditional drift of Z_N, under exact max-X Farmer and exact max-Y Green play and the scaling $s = t/N$, is given by (S1.27)–(S1.29), up to a uniform $o_p(1)$ error. After Farmer exhaustion, the drift is (S1.30).*

Proof. Conditional on the current history, Lemma S1.44 gives iid survivors from $\nu_{x_t, c_{t-1}}$. The conditional X-density at the upper frontier is

$$\frac{\varphi(x_t)\Phi((c_{t-1} - \rho x_t)/\sigma)}{M(x_t, c_{t-1}; \rho)}.$$

Lemma S1.46 gives $n_t = NM(x_t, c_{t-1}; \rho) + O_p(\sqrt{N})$. Hence the expected endpoint spacing of the max-X deletion is

$$\frac{1}{N\,\varphi(x_t)\Phi((c_{t-1} - \rho x_t)/\sigma)} + O_p(N^{-3/2}),$$

which yields $\dot{x}$. The max-Y endpoint spacing gives $\dot{c}$ by symmetry. Farmer's cost is $a(x_t)+o(1)$, and Green's cost is $g_A(x_t, c_{t-1}) + o(1)$, giving the budget drifts. Once Farmer exhausts, the same max-Y calculation applies with x fixed. □

Entrance and launch. For the singular entrance, use tail variables $u = 1-\Phi(x)$, $w = 1-\Phi(c)$, and $q(r) = \Phi^{-1}(1-r)$. The active frontier equations become

$$\dot{u}(s) = F_\rho(u, w) := \frac{1}{\Phi((q(w) - \rho q(u))/\sigma)}, \tag{S1.31}$$

$$\dot{w}(s) = G_\rho(u, w) := \frac{1}{\Phi((q(u) - \rho q(w))/\sigma)}. \tag{S1.32}$$

Lemma S1.48 (Tail-coordinate entrance uniqueness)**.** *For $\rho \in [0, 0.5]$, the system (S1.31)–(S1.32) has a unique entrance solution with $(u(s), w(s)) \to (0, 0)$. It satisfies*

$$\Phi_2(q(u(s)), q(w(s)); \rho) = 1 - 2s, \tag{S1.33}$$

and, for every $0 < \alpha < (1-\rho)/(1+\rho)$,

$$u(s) = s + O_\alpha(s^{1+\alpha}), \qquad w(s) = s + O_\alpha(s^{1+\alpha}). \tag{S1.34}$$

Proof. Interior approximations starting at $\Phi_2(q(r_n), q(r_n); \rho) = 1 - 2s_n$, $s_n \downarrow 0$, yield entrance solutions by Arzela–Ascoli, because on the entrance ratio band the vector field is bounded by $1+C_\alpha s^\alpha$. Along any C^1 solution, differentiating $\Phi_2(q(u), q(w); \rho)$ and using $q'(r) = -1/\varphi(q(r))$

gives derivative -2, hence (S1.33). Since

$$\Phi_2(q(u), q(w); \rho) = 1 - u - w + \Pr(X > q(u), Y > q(w)),$$

with the joint-tail term between 0 and $\min(u, w)$, the mass identity implies $\max\{u(s), w(s)\} \leq 2s$. The lower bound follows from the differential equations themselves: the denominators in (S1.31)–(S1.32) are at most one, so $\dot{u}, \dot{w} \geq 1$; integrating from the entrance, or from the interior approximants and then passing to the limit, gives $u(s), w(s) \geq s$. Therefore $s \leq u(s), w(s) \leq 2s$. On that ratio band, Gaussian quantile asymptotics and Mills' ratio give $F_\rho(u, w), G_\rho(u, w) = 1 + O_\alpha(s^\alpha)$, proving (S1.34). The same estimates give an integrable local Lipschitz bound $C_\alpha s^{\alpha-1}$; Gronwall from the entrance implies uniqueness, and ordinary ODE uniqueness extends it away from zero. $\square$

Lemma S1.49 (Uniform endpoint spacings on the entrance band). *Fix $\rho \in [0, 0.5]$ and $0 < \varepsilon < 1/4$. On histories with $s \leq u, w \leq 2s$, $0 < s \leq \varepsilon$, and $n/N = \Phi_2(q(u), q(w); \rho) + o(1)$, the next exact max-$X$ and max-Y tail-coordinate increments satisfy, uniformly,*

$$\mathbb{E}[\Delta u_N \mid \mathcal{H}] = \frac{1}{N\,\Phi((q(w) - \rho q(u))/\sigma)} + o(N^{-1}), \tag{S1.35}$$

$$\mathbb{E}[\Delta w_N \mid \mathcal{H}] = \frac{1}{N\,\Phi((q(u) - \rho q(w))/\sigma)} + o(N^{-1}), \tag{S1.36}$$

and there exist $C, c > 0$ such that, for all $r \geq 0$,

$$\Pr(N\Delta u_N > r \mid \mathcal{H}) + \Pr(N\Delta w_N > r \mid \mathcal{H}) \leq Ce^{-cr}. \tag{S1.37}$$

Proof. For max-X, the conditional density of the global X-tail coordinate $R = 1 - \Phi(X)$ at its lower endpoint 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 $\rho \leq 0.5$ bound the numerator away from zero and allow standard triangular-array endpoint-order-statistic asymptotics. Thus

$$\mathbb{E}[\Delta u_N \mid \mathcal{H}] = \frac{\Phi_2(q(u), q(w); \rho)}{n\, \Phi((q(w) - \rho q(u))/\sigma)} + o(N^{-1}),$$

which is (S1.35). The strip probability lower bound near the endpoint gives $(1 - dr/N)^n \leq e^{-cr}$ for $z = r/N$, and a fixed-width strip handles large r. The w-coordinate proof is symmetric. □

Proposition S1.50 (Small-share launch in tail coordinates)**.** *Let $\widehat{Z}_N^\rho = (u_N, w_N, b_{F,N}, b_{G,N})$ be the finite-game tail-budget state and let $\widehat{Z}^\rho$ be the entrance solution of (S1.31)–(S1.32) with the corresponding budget equations. For every $\eta > 0$,*

$$\lim_{\varepsilon \downarrow 0} \limsup_{N \to \infty} \Pr\{\|\widehat{Z}_N^\rho(\varepsilon) - \widehat{Z}^\rho(\varepsilon)\| > \eta\} = 0. \tag{S1.38}$$

For each fixed $\varepsilon > 0$, the same conclusion holds for the latent state $Z_N^\rho = (x_N, c_N, b_{F,N}, b_{G,N})$ after applying $x = q(u)$, $c = q(w)$.

Proof. On a small interval $[0, \varepsilon]$, with $\varepsilon < \mu_A/40$, both budgets remain positive with probability tending to one, so exact play and the dynamic survivor law apply. The VC lower-orthant bound and the fact that two labels are removed per round imply $s + o_p(1) \leq u_N, w_N \leq 2s + o_p(1)$ on the entrance window. On this band, Lemma S1.49 gives the conditional drift $(F_\rho, G_\rho, -a(q(u)), -g_A(q(u), q(w)))$, exponential jump tails, and quadratic variation $O(N^{-1})$ for the martingale part. Tightness and subsequential identification therefore give that every limit solves the entrance integral equations. Lemma S1.48 identifies the unique limit, proving fixed-ε convergence; letting $\varepsilon \downarrow 0$ gives (S1.38). For fixed positive ε, $q(\cdot)$ is smooth on the limiting compact tail range, so the continuous mapping theorem transfers convergence to latent coordinates. □

Proposition S1.51 (Stopped fluid convergence). *Let σ_N^G and σ_N^F be the Green and Farmer affordability failure times after share m. For every fixed $\delta > 0$,*

$$\sup_{m \le s \le (\kappa_{\mathrm{low}} - \delta) \wedge \sigma_N^G} \|Z_N(s) - Z(s)\| \xrightarrow{p} 0.$$

The limiting path is the active ODE until $b_F = 0$, followed by the frozen x-frontier post-Farmer ODE.

Proof. Launch from Proposition S1.50 at a small positive share. On compact nonsingular subintervals, Lemma S1.47 gives the predictable drift, and the endpoint-spacing estimates underlying Lemmas S1.49 and S1.55 give exponential tails for the frontier increments and $o_p(1)$ martingale variation over $O(N)$ steps. The stochastic-approximation LLN stated in the technical inputs therefore applies to the latent frontier-budget process. The Farmer-frontier lower envelope in Proposition S1.43 prevents the active path from entering the singular x-tail before Farmer exhaustion. When b_F crosses zero, continuity of b_F and the bounded purchase costs put the finite Farmer-exhaustion window in an $o_p(1)$-neighbourhood of the fluid crossing. After that crossing, x freezes and the same LLN applies to the post-Farmer drift (S1.30). Concatenating the two LLNs gives the displayed stopped convergence. ☐

S1.8.8 Validated $\rho = 0.3$ certificate and affordability bootstrap

The certificate integrates the tail-coordinate ODE on a Picard step-tube. In tail variables $u = 1 - \Phi(x)$, $v = 1 - \Phi(c)$, with $\rho = 0.3$ and $\sigma = \sqrt{1 - \rho^2}$, the validated ODE is

$$\dot{u} = F_1(u, v) := \frac{1}{\Phi((q(v) - \rho q(u))/\sigma)}, \tag{S1.39}$$

$$\dot{v} = F_2(u, v) := \frac{1}{\Phi((q(u) - \rho q(v))/\sigma)}, \tag{S1.40}$$

with the budget drifts

$$\dot b_F = -A(u) := -(10 - 9.9u), \qquad \dot b_G = -G_A(u,v). \tag{S1.41}$$

Here

$$G_A(u,v) := g_A(q(u), q(v)), \tag{S1.42}$$

and, equivalently,

$$G_A(u,v) = 0.1 + 9.9\, \frac{\Phi_2\Big(\frac{q(u)-\rho q(v)}{\sigma}, \frac{\rho q(v)}{\sqrt{2-\rho^2}}; -\frac{\sigma}{\sqrt{2-\rho^2}}\Big)}{\Phi((q(u)-\rho q(v))/\sigma)}. \tag{S1.43}$$

Validated integration specification. The archived scripts use:

(CI1) $m = 0.10$, $\kappa_{\text{low}} = 0.620816$, step size $h = 10^{-4}$, and the two-scale entrance launch at $s_0 = 10^{-5}$.

(CI2) One-sided Abramowitz–Stegun Φ enclosures and bisection inverse normal brackets with outward-rounded interval arithmetic.

(CI3) A Picard tube T_n over each step satisfying

$$B_n + [0,h]\, F(T_n) \subseteq T_n, \tag{S1.44}$$

with a diagonal bootstrap at the entrance.

(CI4) Overlap-weighted accumulation for $c_G = \int_0^{\kappa_{\text{low}}} G_A\, ds$ and $v_{\text{out}} = \int_m^{\kappa_{\text{low}}} e(c(s))\, ds = \int_m^{\kappa_{\text{low}}} (10-9.9v(s))\, ds$.

(CI5) A 96-cell interval quadrature for G_A on each certified tube, with the denominator in

(S1.43) bounded away from zero.

(CI6) A three-phase active/uncertain/post state machine across Farmer exhaustion: active while $\underline{b}_F > 0$, uncertain while $\underline{b}_F \le 0 \le \overline{b}_F$, and post once $\overline{b}_F \le 0$. In the archived run there are 3127 active, 10 uncertain, and 3072 post steps.

Lemma S1.52 (Cost certificate for the $\rho = 0.3$ fluid limit)**.** *Under* (CI1)–(CI6)*,*

$$c_G(\kappa_{\text{low}}) = \int_0^{\kappa_{\text{low}}} G_A(u(s), v(s))\, ds \in [2.353234170771,\ 2.371402296351].$$

Hence

$$\inf_{0 \le s \le \kappa_{\text{low}}} b_G(s) = b_G(\kappa_{\text{low}}) \ge 2.525 - 2.371402296351 = 0.153597703649 > 0.$$

The same certificate encloses Farmer exhaustion by $\tau \le 0.31371$*.*

Proof. The archived interval-arithmetic verifier `rho03_production_certificate.py` certifies the displayed enclosure. At every one of the 6209 steps the script verifies the Picard inclusion (S1.44), a positive lower denominator for the G_A interval quotient, the entrance bootstrap condition, and the active/uncertain/ post state-machine enclosure. The displayed values are those in `rho03_production_certificate.json`. □

Lemma S1.53 (Value certificate for the $\rho = 0.3$ fluid limit)**.** *Under* (CI1)–(CI6)*,*

$$\int_m^{\kappa_{\text{low}}} e(c(s))\, ds = \int_m^{\kappa_{\text{low}}} (10 - 9.9v(s))\, ds \in [2.562783061818,\ 2.585467093762].$$

Combined with $\Theta_{0.3} \le 2.36413$*,*

$$\int_m^{\kappa_{\text{low}}} e(c(s))\, ds - \Theta_{0.3} \ge 0.198653061818 > 0.$$

Proof. The archived interval-arithmetic verifier `rho03_fluid_limit_certificate.py` certifies the enclosure on the same Picard step-tube. Since $10-9.9v$ is decreasing in v, the lower integral enclosure uses the upper endpoint of v on each tube, with post-prefix overlap weights beginning at $m = 0.10$. The displayed interval is the archived JSON value accumulator. $\square$

Proposition S1.54 (Affordability through κ_{low})**.** *Let σ_N^G be the Green affordability failure time in Proposition S1.51. Then*

$$\Pr(\sigma_N^G > \kappa_{\text{low}}) \to 1.$$

Consequently Green plays exact max-Y throughout $[m, \kappa_{\text{low}}]$ with probability tending to one.

Proof. Lemma S1.52 gives a fluid budget floor

$$\beta_0 := \inf_{0 \le s \le \kappa_{\text{low}}} b_G(s) > 0.153597.$$

Fix $\delta > 0$ with $10\delta < \beta_0/4$. Stopped convergence gives

$$\sup_{m \le s \le (\kappa_{\text{low}}-\delta)\wedge\sigma_N^G} |b_G^N(s) - b_G(s)| \to 0 \quad \text{in probability.}$$

If $\sigma_N^G \le \kappa_{\text{low}} - \delta$, then at failure some surviving plot has price above $Nb_G^N(\sigma_N^G)$; since all prices are at most 10, this implies $b_G^N(\sigma_N^G) < 10/N$, contradicting convergence to a value at least β_0. Thus failure before $\kappa_{\text{low}} - \delta$ has vanishing probability. On the terminal strip, the spend over $\delta N + O(1)$ turns is at most $10\delta + o(1)$, so $b_G^N(s) > \beta_0/2$ throughout the strip with probability tending to one; then $Nb_G^N(s) > 10$, and every surviving plot is affordable. Hence no failure occurs before κ_{low}. $\square$

Lemma S1.55 (Selected Green values track the fluid frontier)**.** *On any stopped compact interval*

$[\alpha, \beta] \subset (0, \kappa_{\text{low}}]$ *with exact max-Y play and* $\sup_{\alpha \le s \le \beta} \|Z_N(s) - Z(s)\| \to 0$ *in probability,*

$$\max_{\alpha N \le t \le \beta N} |Y(g_t^{\text{ME}}) - c(t/N)| \xrightarrow{p} 0,$$

and therefore

$$\frac{1}{N} \sum_{t=\lfloor \alpha N \rfloor}^{\lfloor \beta N \rfloor} E(g_t^{\text{ME}}) - \frac{1}{N} \sum_{t=\lfloor \alpha N \rfloor}^{\lfloor \beta N \rfloor} e(c(t/N)) \xrightarrow{p} 0.$$

Proof. On compact interior intervals, the conditional endpoint densities for the max-Y pick are bounded above and below. Lemmas S1.44 and S1.46 give exponential endpoint-spacing tails uniformly over stopped histories. A union bound over $O(N)$ turns yields $\max_t |Y(g_t^{\text{ME}}) - c_N(t/N)| \to 0$ in probability; state convergence gives $c_N \to c$. Uniform continuity of $e(c)$ on the compact range gives the Cesaro value statement. $\square$

Proposition S1.56 (Fluid value convergence up to the certified horizon)**.** *On the* $\rho = 0.3$ ME *path,*

$$\sup_{m \le s \le \kappa_{\text{low}} - \delta} \|Z_N(s) - Z(s)\| \xrightarrow{p} 0 \qquad (\delta > 0),$$

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. Proposition S1.54 makes the stopping in Proposition S1.51 vacuous with probability tending to one, giving compact-subinterval state convergence. Lemma S1.55 and Riemann-sum convergence identify the value sum on $[m, \kappa_{\text{low}} - \delta]$. The terminal strip has at most $\delta N + O(1)$ terms, each bounded by 10, and $\int_{\kappa_{\text{low}} - \delta}^{\kappa_{\text{low}}} e(c(s)) ds \le 10\delta$. Letting $\delta \downarrow 0$ gives the displayed convergence. $\square$

S1.8.9 Static MX threshold and final closure

Lemma S1.57 (Static partition of the MX bottom Farmer block). *Let $B_\phi = \{R \geq r_\phi(0.3)\}$ and*

$$D_{\phi,0.3} = \{A \leq q_F,\ R < r_\phi(0.3)\}, \qquad q_F = Q_A(1-\phi).$$

Under the safe-frontier event of Proposition S1.36, MX*'s Farmer-active Green 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 $D_{\phi,0.3}$, all up to $o_p(N)$ boundary labels. Moreover*

$$B_\phi \mathbin{\dot\cup} D_{\phi,0.3} = \{A \leq q_F\}.$$

Proof. The safe-frontier inequality gives $10/r_\phi(0.3) < q_F$, so $B_\phi \subset \{A \leq q_F\}$ and the population partition follows by splitting $\{A \leq q_F\}$ according to $R \geq r_\phi(0.3)$. The finite-game statement is exactly Proposition S1.36: before Farmer exhaustion, MX buys the initial top-ratio block and Farmer follows the descending A-path; after Farmer stops, the remaining labels below q_F outside the purchased ratio block are the empirical version of $D_{\phi,0.3}$, with only $o_p(N)$ boundary and leftover labels.

We use only the path-independent part of Proposition S1.13: under parity budgets, full leakage, and prices in $[\ell, \nu]$, the number of unclaimed labels at terminal time is at most $\lceil 2\nu/\ell \rceil = 200$, uniformly in N; the subsequent statement there about K_N^{ME} is baseline-specific and is not invoked here. Hence, after Farmer exhaustion, Green purchases all but $O(1)$ labels of the residual region $D_{\phi,0.3}$. This $O(1)$ discrepancy is $o(N)$ in every normalised value sum. $\square$

Proposition S1.58 (MX threshold identity at $\rho = 0.3$). *Under the $\rho = 0.3$ theorem object,*

$$\frac{1}{N}\mathrm{PC}_N^{\mathrm{MX}} \xrightarrow{p} V_G(0.3) = \mathbb{E}_{0.3}[E\,\mathbf{1}\{A \leq q_F\}],$$

and $V_G(0.3) \in [3.26511, 3.26513]$ *by the closed form* (S1.1) *(numerically* $3.265118\ldots$*).*

Proof. By Lemma S1.57, MX's Green purchases are, up to $o_p(N)$ labels, the empirical disjoint union of B_ϕ and $D_{\phi,0.3}$, whose population union is $\{A \le q_F\}$; the thresholded LLN for the bounded weight E gives the limit, and the $O(1)$ terminal leftover (Proposition S1.13) is $o(1)$ when normalised. The value is the $\rho = 0.3$ instance of Lemma S1.3; enclosing the single bivariate-normal CDF gives $V_G(0.3) \in [3.26511, 3.26513]$, so the downstream constant $\Theta_{0.3} \le 2.36413$ of Corollary S1.60 is unchanged. □

Remark S1.59 (The certificate bounds a single shielding deficit)*.* The shielding identity explains the remaining certificate. Since the MX side has the closed-form limit $V_G(0.3)$, any subsequential limit of the payoff gap equals $V_F(0.3) - \mathcal{E}_F^\star(0.3)$, where $\mathcal{E}_F^\star(0.3)$ is the corresponding Farmer ecological capture along the value-first path (Lemma S1.5). The validated fluid certificate of Section S1.8.8 lower-bounds Green's value-first payoff—equivalently, upper-bounds this Farmer capture—by enough to make the gap positive, which is the content of Corollary S1.61.

Corollary S1.60 (Direct theorem reduction via the fluid value path)**.** *Let*

$$\Theta_{0.3} := \mathbb{E}[E\,\mathbf{1}\{A \le q_F\}] - \underline{M}_E(0.10) \le 3.26513 - 0.901 = 2.36413.$$

If

$$\int_m^{\kappa_{\mathrm{low}}} e(c(s))\,ds > \Theta_{0.3},$$

then

$$\liminf_{N\to\infty} \frac{1}{N}\mathbb{E}[\mathrm{PC}_N^{\mathrm{ME}} - \mathrm{PC}_N^{\mathrm{MX}}] > 0.$$

Proof. The prefix lower envelope gives $N^{-1}\sum_{t\le mN} E(g_t^{\mathrm{ME}}) \ge \underline{M}_E(0.10) + o_p(1)$. Proposition S1.56 gives the outside-prefix ME value through κ_{low}, and additional ME purchases are nonnegative. Proposition S1.58 gives the total MX threshold limit. Therefore, in probability,

$$\frac{1}{N}(\mathrm{PC}_N^{\mathrm{ME}} - \mathrm{PC}_N^{\mathrm{MX}}) \geq \int_m^{\kappa_{\mathrm{low}}} e(c(s))\, ds - \Theta_{0.3} + o_p(1).$$

The normalised payoffs are uniformly bounded by 10, so expectation liminf preserves the same lower bound. □

Corollary S1.61 (Validated closure of the $\rho = 0.3$ extension)**.** *With* $m = 0.10$ *and* $\kappa_{\mathrm{low}} = 0.620816$*,*

$$\liminf_{N\to\infty} \frac{1}{N}\mathbb{E}[\mathrm{PC}_N^{\mathrm{ME}} - \mathrm{PC}_N^{\mathrm{MX}}] \geq 2.562783061818 - 2.36413 = 0.198653061818 > 0.$$

Proof. Proposition S1.54 gives affordability through κ_{low}. Proposition S1.56 identifies the ME outside-prefix value with the fluid integral. Lemma S1.53 gives the lower bound 2.562783061818, while Proposition S1.58 and $\underline{M}_E(0.10) = 0.901$ give $\Theta_{0.3} \leq 2.36413$. Apply Corollary S1.60. □

Remark S1.62 (Certificate archival)*.* The scripts `rho03_production_certificate.py` and `rho03_fluid_limit` with their human-readable and JSON reports, are archived on Zenodo (DOI: 10.5281/zenodo.19598799). The JSON reports contain the shared-grid accumulator intervals

$$c_G \in [2.353234170771,\ 2.371402296351], \qquad v_{\mathrm{out}} \in [2.562783061818,\ 2.585467093762].$$

S1.8.10 Standard probabilistic inputs

The proof uses three standard finite-to-fluid tools: uniform empirical convergence for bounded VC-threshold classes, deterministic replacement of $o_p(N)$-equivalent stopping times in bounded normalised sums, and stochastic-approximation fluid convergence for endpoint-order-statistic Markov chains on compact nonsingular intervals. These tools are stated in the baseline proof section and invoked again in the $\rho = 0.3$ extension.

The project-specific work for $\rho = 0.3$ is the dynamic survivor law, the frontier/value-path construction, and the validated affordability certificate in Sections S1.8.5–S1.8.8.

S1.9 Ecological conditions for the knapsack reversal

The asymptotic theorems require sufficient ecological heterogeneity. To characterise this requirement, define the *ecological selection premium*

$$\mathrm{ESP} := \frac{\bar{E}_{\text{top-}e} - \bar{E}_{\text{top-}r}}{\bar{E}},$$

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 ESP $\approx$ 12.5%, while the current proof architecture requires ESP $\approx$ 15%. At the baseline ($E \sim U[0.1, 10]$), 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.

S1.10 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 Rival-Blind 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 S1.63 (Decay of κ)**.** *Under the baseline assumptions:*

(a) $N\,\kappa(N) \xrightarrow{p} \dfrac{20}{5.05} \approx 3.9604$ *as* $N \to \infty$*, where* $\kappa(N) = \max_i a_i/B_g$*.*

(b) For any $c_0 > 20/5.05$*,*

$$P\left(\kappa > \frac{c_0}{N}\right) \;\le\; \exp\left(-\frac{2N\big(2.525 - 10/c_0\big)^2}{24.5025}\right).$$

Proof. Part (a): $\max_i a_i \to 10$ 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 $\max_i a_i \le 10$ 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

$$P(B_g/N < 10/c_0) \;\le\; \exp\left(-\frac{2N\big(2.525 - 10/c_0\big)^2}{24.5025}\right),$$

which establishes part (b). □

Table S1.1 reports the approximate budget concentration ratio $\mathbb{E}[\max_i A_i]/\mathbb{E}[B_g]$ for the grid sizes used in this paper.

Table S1.1: Approximate budget concentration ratio $\mathbb{E}[\max_i A_i]/\mathbb{E}[B_g]$

Grid	N	$\mathbb{E}[\max_i A_i]/\mathbb{E}[B_g]$ (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 S1.64 (What Lemma S1.63 does and does not imply). Lemma S1.63 does not prove the asymptotic dominance theorem; it only shows that the tight-budget mechanism behind the finite counterexamples vanishes under the benchmark sampling scheme. This explains why 2×2 Budget-World pathologies need not persist on the 10×10 baseline or on the Bolivia board. The small-grid exercises below should therefore be read as diagnostics of possible finite-instance reversal mechanisms, not as evidence against the asymptotic theorem.

S1.11 Tiny–grid backward-induction exercises in Budget World

Table S1.2 reports exact backward-induction benchmarks for Budget World on 2×2 and 3×3 grids. The exercise illustrates the tight-budget mechanism behind Proposition 4.3: ratio-greedy can beat value-first on very small boards, but the ranking reverses by 3×3 in these simulations.

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 the mean PC gap over Max Efficiency is positive for every listed ρ, while the majority win rate exceeds one half only for $\rho \geq 0$. 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.

Table S1.2: 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 ± 3.84)	443/2000 (0.63 ± 1.78)	1483/2000 (9.49 ± 8.35)	1892/2000 (18.49 ± 9.09)	271/2000 (−0.86 ± 3.82)
	-0.4	672/2000 (2.58 ± 4.59)	616/2000 (1.17 ± 2.58)	1435/2000 (8.11 ± 7.59)	1813/2000 (15.50 ± 9.34)	361/2000 (−1.42 ± 5.05)
	0.0	857/2000 (2.86 ± 4.62)	746/2000 (1.48 ± 2.96)	1413/2000 (6.76 ± 6.85)	1727/2000 (12.72 ± 9.29)	418/2000 (−1.38 ± 5.36)
	$+0.4$	1009/2000 (3.05 ± 4.38)	869/2000 (1.90 ± 3.29)	1379/2000 (5.58 ± 6.00)	1587/2000 (9.22 ± 8.17)	510/2000 (−1.16 ± 5.39)
	$+0.8$	1134/2000 (3.02 ± 4.04)	1044/2000 (2.18 ± 3.24)	1308/2000 (4.04 ± 4.65)	1417/2000 (5.40 ± 5.66)	638/2000 (−0.85 ± 4.99)
3×3	-0.8	1304/2000 (5.52 ± 6.96)	1489/2000 (6.81 ± 7.68)	1831/2000 (17.73 ± 13.34)	2000/2000 (39.15 ± 13.53)	712/2000 (1.29 ± 6.12)
	-0.4	1410/2000 (5.63 ± 6.43)	1630/2000 (7.10 ± 6.85)	1810/2000 (15.14 ± 11.72)	1998/2000 (33.65 ± 12.92)	897/2000 (1.46 ± 6.91)
	0.0	1443/2000 (5.65 ± 6.28)	1690/2000 (7.44 ± 6.69)	1780/2000 (12.72 ± 10.05)	1982/2000 (27.84 ± 12.03)	1016/2000 (1.79 ± 7.37)
	$+0.4$	1511/2000 (5.55 ± 5.79)	1756/2000 (7.35 ± 6.07)	1766/2000 (10.27 ± 8.43)	1957/2000 (20.49 ± 10.99)	1119/2000 (1.80 ± 7.17)
	$+0.8$	1578/2000 (4.43 ± 4.43)	1781/2000 (6.50 ± 5.22)	1753/2000 (6.81 ± 5.74)	1917/2000 (10.62 ± 6.94)	1197/2000 (2.07 ± 6.26)

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

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.

S2 SPE Diagnostics: full battery

This appendix reports computational subgame-perfect-equilibrium (SPE) diagnostics for both Claims World and Budget World. In each world we ask how our fixed heuristic priority rules perform when the rival Farmer plays a subgame-perfect best response to Green's committed rule, and how those fixed rules compare—in both realised play and payoff—to a subgame-perfect Green conservationist in the full joint SPE against a best-responding Farmer.

The method is common to both worlds: we compute best responses by exact backward induction on the residual board, and run two exercises. In the first, Green commits to a fixed priority rule $Y_\lambda(x) = E_x A_x^{-\lambda}$ and the Farmer best-responds to that commitment—a Stackelberg analysis with a strategic Farmer, not a two-sided SPE. In the second, both players best-respond, yielding the joint SPE, which serves as the benchmark against which fixed-rule performance is measured. We compare each fixed rule to joint-SPE Green along two dimensions: *behavioural resemblance*, using all-action similarity between the realised Green claim sequences, and *commitment performance*, the payoff a fixed rule earns when the Farmer best-responds to it, reported as a share of joint-SPE Green's payoff. Across both worlds the central question is whether the knapsack reversal—established against a mechanical rival in the main theorems—survives when the rival is fully strategic, and how closely the simple priority rules approximate joint-SPE play.

What differs between the worlds is the feasibility structure and the board ensembles. Claims World is pure rivalry, with no prices or budgets, so any remaining plot is always available; Budget World adds prices $p_x = A_x$ and budgets, making feasibility affordability-constrained, and uses a different board ensemble, priority coverage, and sample sizes. We give the Claims-World battery in Section S2.1 and the Budget-World setup in Section S2.8. These diagnostics are evidence about best-response behaviour and priority-rule commitment; they are not part of the formal Budget-World theorem chain in Sections S1.5–S1.8.

S2.1 Claims-World battery design

The Claims-World battery samples boards from 17 distribution classes: independent uniform; Beta$(2,5)$ and Beta$(5,2)$ marginals; compressed-E and very-compressed-E designs; Gaussian-copula rank designs at moderate and high positive and negative dependence; concordant and anti-concordant rank designs; concordant and anti-concordant mixture designs; lognormal-by-rank; random beta marginals; random log-box designs; and integer-lattice stress cases. Sample sizes are $N \in \{4, 6, 8, 10, 12, 14, 16\}$, with 2,000 boards per (N, DGP) cell. The Green priority grid is

$$\lambda \in \{-1.25, -1.1, -1, -0.9, -0.8, -0.7, -0.6, -0.5, -0.25, 0, 0.25, 0.5, 0.75, 1\},$$

supplemented with the A-only Block Farmers rule. Thus there are 238,000 unique boards, each evaluated under each fixed Green rule, giving approximately 3.6 million board–rule evaluations.

Joint-SPE benchmarks are computed by full backward induction on the same boards. Action-match metrics compare the realised Green claim sequence under each fixed rule with the realised Green claim sequence under joint SPE. Payoff-share metrics report each fixed rule's Green payoff divided by joint-SPE Green's payoff. Two features are specific to this battery: the priority grid spans the value-led pre-emption region $\lambda \in [-1, 0]$, where the finite Claims theorem (Theorem 3.1) does not apply, so the diagnostics test whether the priority-tilt monotonicity persists both below $\lambda = 0$ and against a strategic rather than a mechanical Farmer.

S2.2 Claims-World headline λ curve

Table S2.1 reports the headline λ curve, pooling across N and distribution. Mean payoff share is monotone weakly decreasing on $[0, 1]$ (no observed pathwise violations across all 238,000 boards); peaks at or very near $\lambda = -1$ (Hot Spot), with all $\lambda \in [-1.25, -0.9]$ within 0.2 percentage points of the peak; and falls sharply as λ becomes positive. Action-match metrics peak at $\lambda = 0$ (value-first) and decline in both directions, dissociating cleanly from payoff-match metrics.

Table S2.1: Claims World: headline λ curve (pooled N and distribution).

Rule	Mean payoff share	Median payoff share	First-action match	Sequence match
$\lambda = -1.25$	1.004	1.000	0.463	0.205
$\lambda = -1.1$	1.006	1.000	0.477	0.213
Hot Spot ($\lambda = -1$)	**1.006**	1.000	0.498	0.227
$\lambda = -0.9$	1.004	1.000	0.516	0.240
$\lambda = -0.8$	1.002	1.000	0.529	0.248
$\lambda = -0.7$	0.998	1.000	0.545	0.260
$\lambda = -0.6$	0.994	1.000	0.558	0.269
$\lambda = -0.5$	0.987	1.000	0.576	0.286
$\lambda = -0.25$	0.968	1.000	0.629	0.343
Value-first ($\lambda = 0$)	0.945	0.997	**0.725**	**0.482**
$\lambda = 0.25$	0.924	0.971	0.494	0.315
$\lambda = 0.5$	0.904	0.951	0.392	0.238
$\lambda = 0.75$	0.889	0.931	0.327	0.186
Ratio-greedy ($\lambda = 1$)	0.860	0.900	0.261	0.153
Block Farmers (A-only)	0.872	0.886	0.213	0.116

Notes. 238,000 boards per row, pooled over $N \in \{4, 6, 8, 10, 12, 14, 16\}$ and 17 distribution classes. Payoff share is fixed-rule Green payoff divided by joint-SPE Green payoff. First-action match: fraction of boards on which the fixed-rule Green's first claim equals joint-SPE Green's first claim on the realised path. Sequence match: fraction of boards on which the fixed-rule Green's full claim sequence equals joint-SPE Green's full claim sequence. Block Farmers is reported off the λ curve because its priority $Y(x) = A_x$ does not lie in the $E_x A_x^{-\lambda}$ family.

S2.3 SPE-best-response divergence from max-A (Claims World)

The proved Theorem 3.1 concerns the max-A Farmer policy. To show the SPE diagnostics are not parasitic on max-A behaviour, we record the share of boards on which Farmer's subgame-perfect best response diverges from max-A at any node along the realised play path. The divergence rate rises from approximately 21% at $N = 4$ to approximately 85% at $N = 16$, with a decelerating trajectory consistent with an asymptotic limit below 100%. The priority-tilt monotonicity continues to hold without observed violation across this divergence range. The result therefore provides computational evidence for the same monotonicity pattern under

SPE-best-response Farmer behaviour; Appendix S3 closes the gap formally, proving the same monotonicity against the SPE-best-responding Farmer rather than only the mechanical max-A policy.

S2.4 Claims-World joint-SPE comparison: behavioural resemblance versus commitment performance

Table S2.2 separates behavioural resemblance from commitment performance. Value-first is closest to joint-SPE Green on all behavioural similarity metrics. Hot Spot is less mimetic but earns the highest mean payoff as a fixed commitment rule against an SPE-best-responding Farmer. Ratio-greedy and Block Farmers fail for different reasons: the former under-pre-empts contested value, while the latter over-pre-empts without ecological discipline.

Table S2.2: Claims World: selected fixed Green priority rules compared with joint-SPE Green (full all-action battery).

Rule	Mean payoff share	Mean payoff (raw)	First-action match	Full-sequence match	Prefix match	Position-wise match	Set overlap (Jaccard)
Joint SPE	1.000	27.89	1.000	1.000	1.000	1.000	1.000 (1.000)
Hot Spot (EA)	**1.006**	27.93	0.498	0.227	0.301	0.464	0.852 (0.775)
Value-first (E)	0.945	26.22	**0.725**	**0.482**	**0.554**	**0.624**	**0.886** (**0.817**)
Ratio-greedy (E/A)	0.860	24.07	0.261	0.153	0.182	0.290	0.792 (0.687)
Block Farmers (A-only)	0.872	23.80	0.213	0.116	0.140	0.316	0.604 (0.461)

Notes: 238,000 boards per row, pooled over $N \in \{4, 6, 8, 10, 12, 14, 16\}$ and 17 distribution classes. Payoff share is fixed-rule G divided by joint-SPE Green payoff on the same board. First-action match: fraction of boards on which the fixed-rule Green equals joint-SPE Green's first claim. Full-sequence match: fraction on which the entire claim sequence coincides. Prefix ma longest common prefix length divided by joint-SPE Green sequence length. Position-wise match: average across moves of the in the fixed-rule and joint-SPE Green claim the same plot at the same move index. Set overlap: average Jaccard index of final cl reported in parentheses; the headline set-overlap figure is the average $|S_{\text{rule}} \cap S_{\text{SPE}}|/|S_{\text{SPE}}|$. Bold indicates the within-column l fixed rules.

Figure S2.1 visualises the same dissociation: each fixed rule is plotted by its behavioural resemblance to joint-SPE Green against its mean payoff share. Value-first lies furthest right on behavioural resemblance; Hot Spot lies highest on payoff.

The behavioural resemblance ordering is robust across N and distribution class. As N grows, exact sequence match naturally falls because the realised path lengthens and there are more opportunities to diverge: at $N = 16$, value-first's full-sequence match falls to 0.213 and Hot

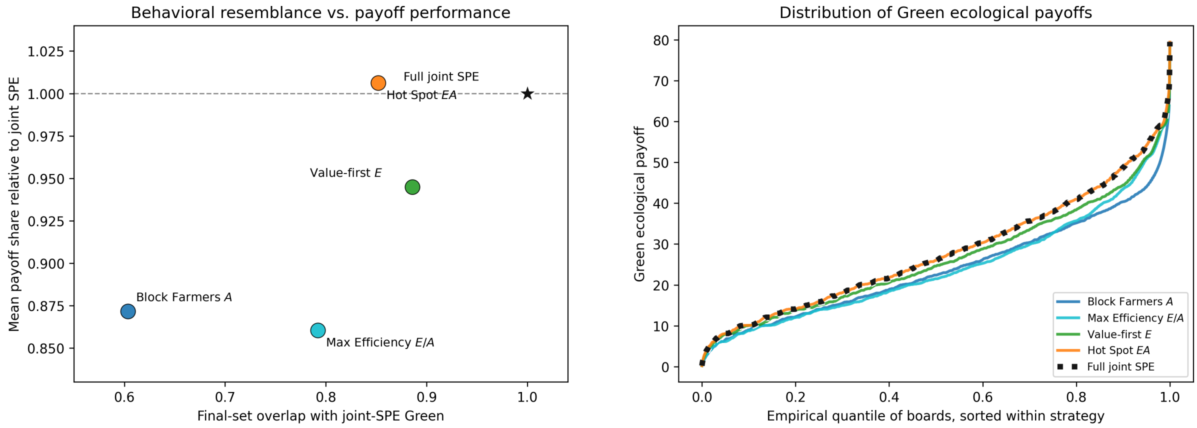

(a) Behavioural resemblance versus payoff **(b)** Payoff distributions

Figure S2.1: Fixed Green priority rules compared with full joint SPE in Claims World. Panel (a) plots, for each fixed rule, final-set overlap with joint-SPE Green against mean payoff share relative to joint SPE. Panel (b) shows empirical payoff quantiles, with boards sorted separately within each strategy.

Notes. 238,000 boards: $N \in \{4, 6, 8, 10, 12, 14, 16\}$, 17 data-generating processes, 2,000 boards per cell. Farmer best-responds by backward induction to each fixed Green rule. Full joint SPE is drawn as a black dotted benchmark. Full battery in Appendix S2.

Spot's to 0.061. But the set overlap ranking remains stable: at $N = 16$ value-first's set overlap is 0.863 and Hot Spot's is 0.840. Set overlap is therefore the more informative behavioural metric at larger board sizes, and the most economically meaningful interpretation of "which fixed rule selects portfolios most like joint-SPE Green's".

S2.5 Per-distribution variation (Claims World)

The headline pattern is robust across distribution classes with two notable structural variations. On anti-concordant rank distributions (where A ranks reverse E ranks), the optimal λ shifts toward $\lambda > 0$: the rival's max-A targets are precisely the low-E plots, so threat-pre-emption is counterproductive and ratio-greedy approaches optimality. On near-collinear distributions (concordant rank, very compressed E), payoff shares cluster near 1 across all λ as the priority rules select nearly identical portfolios. The interior of the distribution space — moderate positive correlation, lognormal-by-rank, mild negative correlation, random log-box, and the uniform baseline — exhibits the cleanest pre-emption pattern with a broad payoff peak in the value-led pre-emption region $\lambda \in [-1, 0]$. Full per-distribution tables are available in the replication package.

S2.6 Per-N trajectory (Claims World)

The mean payoff share by N is approximately monotone in N for each rule, with the value-first share falling from ≈ 0.97 at $N = 4$ to ≈ 0.93 at $N = 16$ and the Hot Spot share remaining close to 1.00 throughout, settling near 1.004 at $N = 16$. Ratio-greedy and Block Farmers shares fall steadily with N to ≈ 0.83 and ≈ 0.84 respectively at $N = 16$. The order-based behavioural similarity metrics (full-sequence match, position-wise match) decline more steeply with N as the realised path lengthens, but the rule-by-rule ordering on these metrics is stable across N: value-first highest, Hot Spot intermediate, ratio-greedy and Block Farmers low. The set-based behavioural similarity metrics (set overlap, Jaccard) are markedly more stable in N: value-first's set overlap holds at 0.86–0.89 across all N, and Hot Spot's at 0.84–0.86. This is the basis for treating set overlap as the most informative behavioural metric at larger board sizes.

S2.7 Methodological reconciliation (Claims World): SPE-best-response Farmer versus heuristic Rival-Aware pre-empting Farmer

The Claims-World simulation results in Section 5.2 feature a heuristic Rival-Aware Farmer with active pre-emption of Green's likely picks. Under that Farmer, value-first outperforms Hot Spot. Under SPE-best-response Farmer, Hot Spot outperforms value-first. The two findings are not in conflict; they reflect different rival behavioural assumptions. SPE-best-response Farmer maximises rival surplus given fixed Green strategy and is therefore the canonical strategic benchmark for procurement-theoretic results. The heuristic Rival-Aware Farmer adds a deny-Green component that goes beyond rival surplus maximisation and tests Green strategies against more aggressively adversarial behaviour. The two together establish that the priority-tilt monotonicity is robust to a range of Farmer behavioural assumptions. The inner-family ranking of value-first vs. Hot Spot, however, is sensitive to that assumption; commitment to Hot Spot is the empirically strongest fixed rule when Farmer is well approximated by SPE best response, while value-first is the safer fallback when the rival's aggressiveness is uncertain.

S2.8 Budget-World SPE diagnostics

This subsection reports the Budget-World counterpart to the Claims-World SPE battery above. The maintained Budget-World protocol is the full-leakage benchmark: prices equal Farmer values, $p_x = A_x$; initial budgets split total price mass 50/50; Green commits to the exact multiplicative order

$$Y_\lambda(x) = E_x A_x^{-\lambda}, \qquad \lambda \in [0, 1],$$

affordability-filtered at each Green move; Farmer moves first and plays a subgame-perfect best response maximising its own terminal $\sum A$ subject to its budget. Thus $\lambda = 0$ is value-first (ME) and $\lambda = 1$ is ratio-greedy (MX). The joint-SPE benchmark below is the non-commitment Budget game in which both players optimise by backward induction at their own nodes; it is not the Stackelberg-optimal commitment policy.

This constant-return objective is the linear boundary of the developer family described in the main text (Figure 6); a fixed tie-breaker resolves the indifference among budget-exhausting portfolios that it leaves. The two regularised objectives that pin down a unique strategic developer, power curvature and a fixed project cost, are summarised in Tables S2.6 and S2.7. Relative to the diagnostics below they sharpen the reversal, raising value-first's efficiency from about nine tenths to 0.95–0.97 and making ratio-greedy wins rare, while leaving its direction and boundaries unchanged.

Validation: the Budget game nests the Claims-SPE game. The Budget protocol contains the Claims-SPE game of Appendix S3 as its non-binding limit. When budgets are slack enough that the affordability filter never removes a feasible plot—the large-budget limit—Green's committed play and Farmer's subgame-perfect best response coincide exactly with the Claims-SPE game, so the SPE-follower monotonicity theorem (Theorem 3.2) applies verbatim and $G_N^{\mathrm{BR}}(\lambda)$ is weakly decreasing on $[0,1]$. Two consequences follow. First, any monotonicity failure in Budget World must originate in the binding-budget regime; the counterexample below is consistent with this, requiring a tight 50/50 split. Second, this furnishes a solver check: in the slack-budget regime the Budget backward-induction solver reproduces the Claims-SPE solver's realised sets, payoffs, and zero monotonicity violations, agreeing in precisely the limit where the two games coincide. Violations appear only as budgets tighten.

Finite monotonicity fails once budgets bind. The Claims-World SPE theorem in Appendix S3 rests on a fixed-order reduction: for a committed Green order, Farmer-realisable terminal sets are the bases of a prefix-laminar matroid. Budget World breaks exactly that structure. Affordability constraints make feasible continuation sets history-dependent: an item can be feasible or infeasible depending on prior budget expenditure, and Green may skip unaffordable high-priority items. The single-prefix exchange argument that proves Claims monotonicity therefore has no analogue in the finite Budget game.

The obstruction is not only formal. Strong finite monotonicity of $G_N^{\mathrm{BR}}(\lambda)$ is false in Budget World, even on continuous boards with no Farmer tie sensitivity. One four-plot counterexample is

	1	2	3	4
A_x	6.6945	8.3998	7.3569	4.7397
E_x	2.5048	5.4719	4.9747	3.6902

with 50/50 budgets. Along the committed Y_λ curve, Green's purchased conservation rises by approximately 0.4973 at an interior step as λ increases, so the Claims-style statement "G_N^{BR} is weakly decreasing on $[0,1]$" cannot be extended to finite Budget World. The counterexample is not tie-break driven. It is a genuine budget effect: the ordering change interacts with residual affordability, not with non-unique Farmer payoffs.

The correct Budget question is therefore not finite pathwise monotonicity but endpoint and benchmark performance: how value-first compares with ratio-greedy against an SPE-best-responding Farmer, and how both simple rules compare with full non-commitment joint SPE.

Fixed-rule endpoint performance and joint-SPE efficiency. Table S2.3 summarises the main $N = 16$ joint-SPE benchmark. The ordinary heterogeneous cells are the continuous uniform, positive-copula, beta, and continuous-lognormal designs. In those cells, value-first captures 91.4% of the joint-SPE Green payoff on average, while ratio-greedy captures 89.5%. The mean fixed-rule endpoint gap $G^{\mathrm{ME}} - G^{\mathrm{MX}}$ is 0.899 Green payoff units, with a 95% confidence interval $[0.750, 1.049]$. Adding the mild anti-concordant cells $\rho = -0.6$ and $\rho = -0.7$ leaves the mean gap positive. By contrast, compressed ecological values and the full strong anti-concordance sweep are scope boundaries: the mean gap is negative or statistically indistinguishable from zero, even though weak non-loss remains common.

Table S2.3: Budget World: fixed rules compared with joint SPE at $N = 16$.

Cell group	Boards	G^{ME}/G^J	G^{MX}/G^J	$G^{\mathrm{ME}} - G^{\mathrm{MX}}$	95% CI	$\Pr(G^{\mathrm{ME}} \geq G^{\mathrm{MX}})$
Ordinary heterogeneous	2,000	0.914	0.895	0.899	[0.750, 1.049]	0.814
Ordinary $+\rho = -0.6, -0.7$	3,000	0.907	0.894	0.621	[0.495, 0.747]	0.832
Negative-ρ sweep	2,500	0.895	0.896	-0.014	[-0.181, 0.153]	0.876
Compressed E	500	0.896	0.914	-0.861	[-1.176, -0.546]	0.694
All cells	5,000	0.903	0.897	0.267	[0.158, 0.375]	0.833

Notes: G^J is Green's payoff in the non-commitment joint SPE of the same Budget game. G^{ME} and G^{MX} are the fixed-rule payoffs against a Farmer SPE best response to the committed rule. "Ordinary heterogeneous" pools the continuous uniform, positive-copula, beta, and continuous-lognormal cells. The negative-ρ sweep pools $\rho \in \{-0.6, -0.7, -0.8, -0.9, -0.95\}$. All designs are continuous and tie-audited; fixed-rule endpoint comparisons are tie-stable in these cells.

The table should be read as an efficiency benchmark, not as a claim that joint SPE is a formal upper bound. A fixed commitment rule can outperform the non-commitment SPE on some boards because commitment has strategic value. In Budget World, however, joint SPE is practically ceiling-like in expectation: value-first is below the joint-SPE payoff on average but reaches about nine tenths of it, and it does so while outperforming the textbook ratio rule in the ordinary heterogeneous cells.

Convergence of the efficiency ratio. The stable asymptotic object in the Budget-SPE computations is the intensive ratio G^{ME}/G^J, not the raw endpoint gap. Table S2.4 reports the convergence pattern over $N = 4$ through 16 for the ten-cell continuous benchmark used in Table S2.3. The value-first efficiency ratio falls from the small-board optimism of 0.989 at $N = 4$ and stabilises around 0.90 by $N \approx 12$–16. Ratio-greedy also declines from its small-board value and stabilises below value-first over the same range.

Table S2.4: Budget World: convergence of fixed-rule efficiency relative to joint SPE.

N	Boards	G^{ME}/G^J	G^{MX}/G^J	$G^{\mathrm{ME}} - G^{\mathrm{MX}}$	$\Pr(G^{\mathrm{ME}} \geq G^{\mathrm{MX}})$
4	5,000	0.989	0.948	0.326	0.978
6	5,000	0.964	0.929	0.541	0.956
8	5,000	0.941	0.916	0.561	0.927
10	5,000	0.923	0.906	0.504	0.890
12	5,000	0.911	0.898	0.473	0.866
14	5,000	0.906	0.898	0.353	0.841
16	5,000	0.903	0.897	0.267	0.833

Notes: pooled over ten continuous Budget-SPE cells: uniform, positive-copula, beta, compressed-E, continuous-lognormal, and $\rho \in \{-0.6, -0.7, -0.8, -0.9, -0.95\}$. The endpoint gap is reported in raw Green payoff units. Its normalisation by total ecological mass declines with N; the ratio to joint SPE is the stable intensive object.

This distinction matters for interpretation. The fixed-rule endpoint gap is positive and persistent in ordinary heterogeneous environments, but the data do not support a Bronze-style positive ecological-mass intensity against an SPE Farmer. The supported benchmark claim is instead that value-first attains a stable fraction of strategic non-commitment play.

Correlation and heterogeneity scope. The correlation frontier refines the scope conditions. Positive concordance between A and E strengthens the value-first edge: high-E plots become more likely to be high-A plots, hence more likely to be contested by the Farmer and more costly for ratio-greedy to pursue. In the symmetric positive copula sweep over $N = 4, 6, 8, 10, 12, 14$, the mean gap $G^{\mathrm{ME}} - G^{\mathrm{MX}}$ rises from 0.862 at $\rho = 0$ to 1.481 at $\rho = 0.95$; value-first's efficiency relative to joint SPE remains in the low-to-mid 90% range, while ratio-greedy falls toward the high 80% range. This monotone profile is consistent with Figure 5 in the main text, which reports the normalised Purchased Conservation gap against the heuristic Rival-Blind Farmer: both keep the value-first advantage positive and increasing across the empirically relevant range, and they differ only in the far positive tail, where the heuristic-rival gap is hump-shaped while the strategic-rival gap continues to rise. The advantage therefore does not weaken against a more sophisticated rival.

Table S2.5: Budget World: positive-correlation sweep, $N = 4$–14.

ρ	G^{ME}/G^{J}	G^{MX}/G^{J}	$G^{\mathrm{ME}} - G^{\mathrm{MX}}$	$\Pr(G^{\mathrm{ME}} \geq G^{\mathrm{MX}})$
0.00	0.935	0.898	0.862	0.898
0.30	0.931	0.892	0.884	0.863
0.50	0.932	0.885	1.023	0.859
0.70	0.935	0.876	1.253	0.848
0.90	0.939	0.868	1.407	0.834
0.95	0.943	0.867	1.481	0.836

Notes: 500 boards per (N, ρ) cell, pooled over $N \in \{4, 6, 8, 10, 12, 14\}$. $\rho = 0$ is the independent Gaussian-copula uniform-margin benchmark. The weak non-loss rate includes endpoint ties.

The negative side of the frontier behaves differently. Moderate anti-concordance attenuates the mean gap but leaves the typical-board statement intact. At large N, value-first usually ties or beats ratio-greedy even for strongly negative ρ, but the expectation becomes tail-sensitive: a small left tail of large ratio-greedy wins can overturn the mean. In the fixed-rule $N = 18$ screen, for example, $\rho = -0.95$ has $\Pr(G^{\mathrm{ME}} \geq G^{\mathrm{MX}}) \approx 0.872$ and median gap zero, while the mean gap is negative. Compressed ecological heterogeneity is an independent failure regime: when the dispersion in E is too small, value-first has little selection premium to exploit and the mean reversal collapses.

The left tail as a commitment artifact. The joint-SPE benchmark explains the bad tail. On boards where value-first loses to ratio-greedy, the flexible joint-SPE Green usually performs well. In the ordinary heterogeneous $N = 16$ cells, conditional on $G^{\mathrm{ME}} < G^{\mathrm{MX}}$, the average payoffs are

$$\mathbb{E}[G^J \mid G^{\mathrm{ME}} < G^{\mathrm{MX}}] = 46.56, \qquad \mathbb{E}[G^{\mathrm{ME}} \mid G^{\mathrm{ME}} < G^{\mathrm{MX}}] = 39.43, \qquad \mathbb{E}[G^{\mathrm{MX}} \mid G^{\mathrm{ME}} < G^{\mathrm{MX}}] = 42.69.$$

Thus the boards are not intrinsically hard for Green: a flexible Green escapes the trap. The value-first shortfall is the cost of committing to a rigid priority rule that sometimes overcommits to expensive high-E plots under an adversarial budget. Across all $N = 16$ cells, joint SPE scores about 22% higher than value-first on value-first-loss boards; in the ordinary heterogeneous cells the corresponding figure is about 18%, and rises above 25% in the strong anti-concordance tail. The worst decile of value-first relative-to-ratio boards accounts for about one fifth of the total joint-minus-value-first shortfall, so the commitment cost is partly tail-driven but not only a tail phenomenon.

Hot Spot versus value-first: pre-emption is priced differently across worlds. The Claims diagnostics above show that explicit pre-emption can pay when it is free. Hot Spot is the off-interval member of the same multiplicative family,

$$Y_{-1}(x) = E_x A_x,$$

which tilts value-first toward plots that are both ecologically valuable and highly valued by Farmer. In Claims World this tilt is costless: Green does not spend a budget, so a committed Hot Spot order can beat value-first against the same SPE-best-responding Farmer. Budget World reverses the interpretation. Since price equals A, Hot Spot asks Green to pay extra precisely for the plots Farmer likes. Once threat is priced, explicit pre-emption can overpay for contested plots and value-first becomes the safer simple commitment rule. Table S2.6 compares value-first and Hot Spot against an SPE-best-responding Farmer across the three developer objectives. In the free-claiming Claims battery Hot Spot dominates, beating value-first on all but a handful of boards. Pricing changes this, but by objective: under the linear boundary value-first clearly overtakes Hot Spot, under the fixed cost it does so modestly, and under power curvature the two are statistically tied. In every objective the best rule is neither endpoint but the interior peak, which beats both. The value of explicit pre-emption is thus not eliminated by pricing so much as moderated: full pre-emption loses the commanding advantage it holds under free claiming, while a moderate, value-led tilt remains best.

Table S2.6: Value-first versus Hot Spot against an SPE-best-responding Farmer, by developer objective ($N = 16$, $\rho = 0$, 500 boards).

Developer objective	ME (% SPE)	Hot Spot (% SPE)	Peak ($\lambda^\star$) (% SPE)	Pr(ME $\geq$ HS)	$\mathbb{E}[G^{\text{ME}} - G^{\text{HS}}]$
Linear boundary	88.8	82.9	89.3 (+0.13)	0.858	+3.13
Power curvature	95.2	95.6	96.7 (−0.38)	0.636	−0.10
Fixed cost	96.7	96.2	97.8 (−0.38)	0.704	+0.43

Notes: efficiency is Green payoff as a percentage of the objective-specific non-commitment joint-SPE payoff. ME is value-first ($Y_0 = E$), Hot Spot is $Y_{-1} = EA$; Peak is the interior maximum at tilt $\lambda^\star$. The last two columns compare per-board Green payoffs directly (weak inequality; raw ecological units). For contrast, in the free-claiming Claims battery value-first weakly beats Hot Spot on only 4.1% of boards ($\mathbb{E}[G^{\text{ME}} - G^{\text{HS}}] = -6.41$): Hot Spot dominates when pre-emption is free.

Extending this comparison from the value-first/Hot Spot pair to all four headline rules places them on the priority-tilt axis and makes the world-specific structure explicit (Table S2.7; the corresponding curves are Figure 6 in the main text). Efficiency relative to the non-commitment joint-SPE benchmark has an interior peak in every case, and the peak's location shifts with the world and the developer's objective: it sits at Hot Spot ($\lambda = -1$) in Claims World, near value-first under the linear Budget boundary, and at a moderate pre-emptive interior under the power and fixed-cost objectives. Both naive extremes underperform the interior throughout—ratio-greedy ($\lambda = 1$) on one flank and the A-only Block endpoint ($\lambda \to -\infty$) on the other—with Block the floor in every case (from 0.65 to 0.75 across the Budget objectives), where it spends budget on the most expensive plots. In Claims World the Hot Spot anchor exceeds one, confirming that a committed rule can outperform the non-commitment joint-SPE benchmark; in Budget World every anchor lies below one, so pricing removes that premium.

Table S2.7: Priority-tilt anchors: efficiency relative to the objective-specific non-commitment joint SPE against an SPE-best-responding Farmer ($N = 16$, 500 boards).

	Block (A) $\lambda \to -\infty$	Hot Spot (EA) $\lambda = -1$	Value-first (E) $\lambda = 0$	Ratio (E/A) $\lambda = +1$
Claims World	0.800	1.002	0.878	0.803
Budget: linear	0.654	0.829	0.888	0.873
Budget: power	0.740	0.956	0.952	0.932
Budget: fixed cost	0.753	0.962	0.967	0.942

Notes: each entry is Green's payoff under the committed rule divided by the non-commitment joint-SPE Green payoff on the same board, averaged over 500 boards. The interior peak (not shown) exceeds every anchor: near value-first under the linear boundary and in a moderately pre-emptive region under the power and fixed-cost objectives. The Claims Hot Spot anchor exceeds one, a commitment premium that pricing removes.

Proof status and open theorem target. The Budget-SPE object suggested by the computations is an approximation-ratio or competitive-efficiency statement, not a finite monotonicity theorem. A representative target would be to show that, for a class of continuous heterogeneous distributions excluding compressed E and sufficiently strong anti-concordance,

$$\frac{\mathbb{E}[G_N^{\mathrm{ME}}]}{\mathbb{E}[G_N^{J}]} \longrightarrow q(\mu), \qquad q(\mu) \approx 0.90$$

for the distribution μ governing (A, E). Proving such a statement would require a new fluid or mean-field limit for a two-player alternating budget game: the state would have to track the empirical measure of remaining (A, E) plots, the two budget coordinates, and the turn/skip phase, and then compare the limiting value functional of a committed rule with the limiting joint-SPE value functional. This is a different order of difficulty from the Claims proof. The finite prefix-matroid machinery that proves the Claims SPE-follower theorem does not survive binding budgets, and the joint-SPE benchmark introduces two optimising players rather than one. We therefore report the Budget-SPE results as a diagnostic and benchmark pattern, and leave the approximation-ratio theorem as an open target.

S3 Claims World: SPE-follower priority-tilt monotonicity

This section upgrades the finite-board Claims theorem (Theorem 3.1) from the *mechanical* max-A Farmer to the *subgame-perfect best-responding* Farmer. The object is the Stackelberg-follower game underlying the SPE diagnostics of Appendix S2: Green commits to a fixed priority order, Farmer best-responds by backward induction to maximise its own captured agricultural value $\sum_{x \in F} A_x$, and Green then takes the highest-priority remaining plot at each of its moves. Throughout this section, "SPE-follower" means Farmer's subgame-perfect best response in the finite perfect-information game induced by that fixed committed Green priority rule; it is not a claim about a joint SPE in which Green dynamically optimises. We prove that, on the priority interval $\lambda \in [0, 1]$, Green's realised ecological payoff is weakly decreasing in λ: value-first weakly dominates ratio-greedy against this best-responding Farmer.

The argument is *not* a corollary of Theorem 3.1. In the diagnostic battery (Appendix S2.3) the SPE-follower Farmer diverges from max-A play on a majority of boards, and in those observed divergences it captures strictly more A and leaves Green strictly less E. What survives is not the Farmer's path but the *direction* in which committed-order changes move Farmer's feasible capacity. The proof isolates that direction by collapsing the backward induction to a static maximum-weight basis problem on a laminar matroid, and then reusing the adjacent-crossing architecture of Theorem 3.1.

Setup and conventions

Fix an even board $N = 2K$ with plot set X, agricultural values $A_x > 0$ and ecological values $E_x > 0$. Green commits to a strict priority order $\sigma = (x_1, \ldots, x_{2K})$ (x_1 highest priority). Play is *Farmer-first* and strictly alternating, $F, G, F, G, \ldots$, for K rounds; on each Green move Green claims the highest-σ remaining plot. Farmer's payoff is terminal and additively separable in its final claimed set, $u_F(F) = \sum_{x \in F} A_x$; Farmer is indifferent among histories that generate the same final set, and there is no intermediate, order-dependent, or blocking payoff. Write

$$P_m^\sigma = \{x_1, \ldots, x_m\}$$

for the length-m prefix of σ.

When a priority order is generated by scores $Y_\lambda(x) = E_x A_x^{-\lambda}$, fix once and for all an arbitrary strict exogenous tie-break order τ on X. Let $\sigma(\lambda)$ denote the strict order that sorts plots by decreasing $Y_\lambda(x)$, breaking ties by τ. Thus $G_N^{\mathrm{BR}}(\lambda)$ is a well-defined function on the closed interval $[0, 1]$.

A board is *generic* if: (i) all A_x are distinct; (ii) all E_x are distinct; (iii) for every $\lambda \in (0, 1)$ at most one unordered pair ties in Y_λ, with no endpoint ties; and (iv) *no two distinct subsets of X have equal A-sum*. Conditions (i)–(iii) are those of Theorem 3.1; (iv) is the additional *no-Farmer-payoff-tie* condition needed to make the SPE Farmer set, and hence Green's realised payoff, well-defined. Conditions (i)–(iv) hold almost surely whenever the joint law of $(A_x, E_x)_{x \in X}$ is absolutely continuous with respect to Lebesgue measure on $\mathbb{R}_+^{2N}$. This is a generic-position theorem: integer, categorical, or otherwise tie-heavy boards need not satisfy (iii) or (iv). On such non-generic boards Farmer's best-response set can be a correspondence and Green's realised payoff can depend on an explicit equilibrium selection or tie-breaking convention; the theorem below makes no claim for those cases.

The fixed-order reduction

Lemma S3.1 (Realisability)**.** *A set $F \subseteq X$ with $|F| = K$ is achievable as Farmer's final set, for some Farmer strategy against the committed order σ, if and only if*

$$|F \cap P^{\sigma}_{m}| \leq \left\lceil \tfrac{m}{2} \right\rceil \qquad \textit{for every } m \in \{1, \dots, N\}. \tag{S3.1}$$

Proof. Necessity. Let $G = X \setminus F$ be Green's final set and let $q_1 < \cdots < q_K$ be the σ-positions of its plots. Green claims them in this order: by induction, before Green's t-th move Green owns $x_{q_1}, \dots, x_{q_{t-1}}$, and Green's greedy rule then selects the lowest-position remaining plot, which is x_{q_t} provided every lower-position plot outside G is already removed. The lower-position plots outside G number $q_t - 1 - (t-1) = q_t - t$, and these must have been claimed by Farmer; Farmer has made exactly t moves by Green's t-th move (Farmer first), so $q_t - t \leq t$, i.e. $q_t \leq 2t$. Equivalently $|G \cap P^{\sigma}_{2t}| \geq t$, i.e. $|F \cap P^{\sigma}_{2t}| \leq t$; adding the single plot x_{2t+1} gives the odd-prefix bound, establishing (S3.1).

Sufficiency. Suppose (S3.1) holds. Let $G = X \setminus F$ and let $q_1 < \cdots < q_K$ be the σ-positions of the elements of G. The prefix inequalities at even lengths give $|F \cap P^{\sigma}_{2t}| \leq t$, hence $|G \cap P^{\sigma}_{2t}| \geq t$, which is exactly

$$q_t \leq 2t \qquad \text{for every } t = 1, \dots, K \tag{S3.2}$$

(if $q_t > 2t$ then the first $2t$ positions hold at most $t-1$ elements of G, so at least $t+1$ of F, violating $|F \cap P^{\sigma}_{2t}| \leq t$). Let Farmer play the *earliest-target* strategy: on each move claim the lowest-σ-position currently-unclaimed plot of F. We prove by induction on t that, immediately before Green's t-th move:

(a) Farmer has claimed the t lowest-σ-position elements of F;

(b) Green has claimed $x_{q_1}, \ldots, x_{q_{t-1}}$;

(c) every unclaimed plot with σ-position below q_t has been removed, so Green's greedy rule forces x_{q_t}.

Base $t = 1$. Farmer has made one move and claimed the lowest-position element of F; Green has claimed nothing. Among positions $1, \ldots, q_1 - 1$ every plot lies in F, and there are $q_1 - 1 \le 1$ of them by (S3.2); Farmer's single move has therefore removed all of them, so x_{q_1} is the lowest-position remaining plot and Green is forced to it. *Step* $t - 1 \Rightarrow t$. By hypothesis Green owns $x_{q_1}, \ldots, x_{q_{t-1}}$ and Farmer owns the $t-1$ lowest-position elements of F; Farmer's t-th move claims the next, giving (a). (Green's ownership of $x_{q_{t-1}}$ here is the forced move established by part (c) at stage $t-1$, so immediately before Farmer's t-th move Green owns $x_{q_1}, \ldots, x_{q_{t-1}}$.) Among the first $q_t - 1$ positions exactly $t - 1$ belong to G (namely $x_{q_1}, \ldots, x_{q_{t-1}}$), so the remaining $(q_t - 1) - (t - 1) = q_t - t$ belong to F; by (S3.2), $q_t - t \le t$, so every F-plot below q_t is among the t Farmer now owns and has been claimed. The G-plots below q_t were claimed by Green by (b). Hence every plot below q_t is removed while x_{q_t} remains, giving (c). The induction closes: Green claims exactly G in the order $x_{q_1}, \ldots, x_{q_K}$ and Farmer ends with exactly F, so F is realisable.

Remark on the repair. The decisive step is the count $q_t - t \le t$, not the (false) claim that Farmer's t-th *target* precedes x_{q_t} in σ: Farmer may well target a plot lying after x_{q_t}. What the count guarantees is the weaker and correct fact that all F-plots *below* x_{q_t} have already been cleared, which is what forces Green onto x_{q_t}. □

Corollary S3.2 (Laminar matroid)**.** *The family* $\mathcal{I}_\sigma = \{S \subseteq X : |S \cap P^\sigma_m| \le \lceil m/2 \rceil \ \forall m\}$ *is a matroid: it is the laminar matroid on the chain* $P^\sigma_1 \subset \cdots \subset P^\sigma_N$ *with capacities* $\lceil m/2 \rceil$*. Its bases are exactly the size-*K *realisable Farmer sets of Lemma S3.1.*

Proof. This is the standard laminar-matroid construction specialised to a chain. For completeness, here is the exchange argument. Heredity is immediate. Let $I, J \in \mathcal{I}_\sigma$ with $|I| < |J|$, and choose $x_r \in J \setminus I$ with largest σ-position r. If $I \cup \{x_r\}$ violated independence, some prefix P^σ_m with $m \geq r$ would be saturated by I: $|I \cap P^\sigma_m| = \lceil m/2 \rceil$. Since x_r was chosen as the latest position in $J \setminus I$, every element of $J \setminus P^\sigma_m$ already lies in I. Hence

$$|J| = |J \cap P^\sigma_m| + |J \setminus P^\sigma_m| \leq \lceil m/2 \rceil + |I \setminus P^\sigma_m| = |I|,$$

contradicting $|I| < |J|$. Thus $I \cup \{x_r\} \in \mathcal{I}_\sigma$, proving the exchange axiom. Finally $P^\sigma_N = X$ has capacity K, so bases have size K; Lemma S3.1 identifies exactly these bases with the realisable Farmer sets. □

Lemma S3.3 (SPE follower equals maximum-A basis)**.** *On a generic board, Farmer's subgame-perfect best response to the committed order σ yields the unique maximum-A basis*

$$B_\sigma = \arg\max\Big\{\textstyle\sum_{x \in F} A_x : \ F \ \textit{a basis of} \ \mathcal{I}_\sigma\Big\},$$

computable by the matroid greedy algorithm. Green's realised ecological payoff is

$$G^{\mathrm{BR}}(\sigma) = \sum_{x \in X} E_x - \sum_{x \in B_\sigma} E_x.$$

Proof. The game induced by a fixed committed order σ is finite, deterministic, and of perfect information: at Farmer nodes Farmer chooses one remaining plot, and at Green nodes the transition is fixed by σ. Farmer's continuation value at any Farmer node is therefore the maximum terminal A-payoff over terminal Farmer sets reachable from that history, and the usual finite-horizon backward-induction recursion selects an action attaining this maximum at every Farmer node. Thus the Farmer strategy can be completed off path to a pure subgame-perfect best response.

At the initial history, every basis of $\mathcal{I}_\sigma$ is achievable (Lemma S3.1, sufficiency) and every achieved Farmer set is a basis (necessity, with $|F| = K$). Since Farmer's payoff is terminal and additive, the initial backward-induction terminal set must therefore be the maximum-A basis of $\mathcal{I}_\sigma$. Genericity (iv) makes the maximiser unique, so Green's realised payoff is well-defined. The matroid greedy algorithm computes the maximum-weight basis of a matroid, giving B_σ. $\square$

Remark S3.4. This is the precise sense in which the SPE Farmer is "max-A subject to Green-priority constraints": not max-A along the realised path, but the maximum-A *basis* of the prefix matroid $\mathcal{I}_\sigma$. The path-level divergence from mechanical max-A (Appendix S2.3) is exactly the gap between greedy-available play and the matroid optimum.

The adjacent-crossing exchange

We now show each ratio-direction crossing weakly lowers G^{BR}. The load-bearing step is an exchange lemma; the key structural fact is that the two swapped plots are *symmetric* in the matroid obtained by dropping the one constraint that distinguishes the two orders, so that Farmer's strict A-preference forces a single clean exchange. This is what rules out the "hidden multi-exchange" in the odd-prefix case.

We use one standard matroid fact, the *bijective basis-exchange theorem* of Brualdi Brualdi (1969): for any two bases B, C of a matroid there is a bijection $\phi : B \setminus C \to C \setminus B$ such that $B \setminus \{b\} \cup \{\phi(b)\}$ is a basis for every $b \in B \setminus C$. We use only the one-sided form of this theorem: for the single bijection ϕ, each $B \setminus \{b\} \cup \{\phi(b)\}$ is a basis. We do *not* invoke *base-orderability*, the strictly stronger property requiring a bijection under which $B \setminus \{b\} \cup \{\phi(b)\}$ and $C \setminus \{\phi(b)\} \cup \{b\}$ are *simultaneously* bases, which does not hold for general matroids.

Lemma S3.5 (One-cut basis sensitivity). *Let M_0 be a matroid of rank K on ground set X with weights A_x having no ties across bases, let $L \subseteq X$, and fix an integer r. Write $A(S) = \sum_{x \in S} A_x$, let B_0 be the unique maximum-A basis of M_0, and let B_1 be the unique maximum-A basis among* *bases B with $|B \cap L| \leq r$ (assume one exists). If $|B_0 \cap L| \leq r+1$, then either $B_1 = B_0$, or*

$$B_1 = B_0 \setminus \{z\} \cup \{w\}, \qquad z \in B_0 \cap L, \quad w \in X \setminus L, \quad A_z > A_w.$$

Proof. If $|B_0 \cap L| \leq r$ then B_0 itself satisfies the side constraint, and being M_0-optimal it is optimal in the smaller class, so $B_1 = B_0$ by uniqueness. Suppose $|B_0 \cap L| = r+1$. Let C be any basis with $|C \cap L| \leq r$ and apply Brualdi exchange to B_0, C, giving $\phi : B_0 \setminus C \to C \setminus B_0$ with each $B_0 \setminus \{b\} \cup \{\phi(b)\}$ a basis. Since $|(B_0 \setminus C) \cap L| = |B_0 \cap L| - |B_0 \cap C \cap L|$ exceeds $|(C \setminus B_0) \cap L| = |C \cap L| - |B_0 \cap C \cap L|$, the map ϕ cannot send every element of $(B_0 \setminus C) \cap L$ into L; pick $z \in (B_0 \setminus C) \cap L$ with $w := \phi(z) \notin L$ and set $C' = B_0 \setminus \{z\} \cup \{w\}$. Then C' is a basis, and since $z \in L$, $w \notin L$, $|B_0 \cap L| = r+1$ we have $|C' \cap L| = r$, so C' is in the constrained class. Because B_0 is maximum-A over M_0, every exchange gives $A_b \geq A_{\phi(b)}$, so $A(B_0) - A(C) = \sum_{b \in B_0 \setminus C}(A_b - A_{\phi(b)})$ has all summands nonnegative; in particular $A_z - A_w \leq A(B_0) - A(C)$, whence $A(C') = A(B_0) - A_z + A_w \geq A(C)$. Taking $C = B_1$, the basis C' is in the constrained class with $A(C') \geq A(B_1)$, so $C' = B_1$ by optimality and uniqueness; thus $B_1 = B_0 \setminus \{z\} \cup \{w\}$. Finally $A_z > A_w$. Here $w \notin B_0$ because $w = \phi(z) \in C \setminus B_0$, so $B_0 \setminus \{z\} \cup \{w\}$ is distinct from B_0. If $A_z < A_w$ then this would be a strictly heavier M_0 basis, and if $A_z = A_w$ a distinct equal-weight one, either contradicting the optimality/uniqueness of B_0. □

Lemma S3.6 (Adjacent-crossing exchange)**.** *Let $\sigma^- = P\,i\,j\,T$ and $\sigma^+ = P\,j\,i\,T$ be the committed orders just before and just after a generic interior crossing $\lambda_* \in (0,1)$, with i ahead of j on the left. Then $A_i > A_j$, $E_i > E_j$, $Y_{\lambda_*}(i) = Y_{\lambda_*}(j)$, every plot of P has $Y_{\lambda_*} \geq$ this tied score, and every plot of T has $Y_{\lambda_*} \leq$ it. Writing $B^- = B_{\sigma^-}$, $B^+ = B_{\sigma^+}$,*

$$\sum_{x \in B^+} E_x \;\geq\; \sum_{x \in B^-} E_x, \qquad \textit{equivalently} \qquad G^{\mathrm{BR}}(\sigma^-) \;\geq\; G^{\mathrm{BR}}(\sigma^+).$$

Proof. The crossing facts are as in Theorem 3.1: since i leads left of λ_* and trails right of it,

i is the faster-falling (hence higher-A) plot, $A_i > A_j$; and $d(\lambda) = \log(E_i/E_j) - \lambda \log(A_i/A_j)$ vanishes at $\lambda_* > 0$ with $\log(A_i/A_j) > 0$, so $E_i > E_j$. Plots of P lead the tied pair at λ_* and plots of T trail it.

Step 0: only one constraint differs. Set $p = |P|$. For $m \le p$ both orders have prefix P; for $m \ge p+2$ both have the same prefix $P \cup \{i, j\} \cup$ (first $m - p - 2$ of T). The orders differ only at $m = p + 1$, where σ^- imposes $|S \cap (P \cup \{i\})| \le \lceil (p+1)/2 \rceil$ and σ^+ imposes $|S \cap (P \cup \{j\})| \le \lceil (p+1)/2 \rceil$. Let M_0 be the common matroid with this single constraint dropped. In M_0 the only constraints touching i or j are the prefixes $m \ge p + 2$, which contain *both* and are symmetric under $i \leftrightarrow j$; so i and j are exchangeable in M_0. By genericity condition (iv), any two distinct bases of any matroid considered below, including M_0, have distinct A-weights; hence all maximum-A bases below are unique. Let $c = \lceil p/2 \rceil$ be the cap of the $m = p$ constraint $|S \cap P| \le c$ (which lies in M_0).

Step 1: p even. Then $\lceil (p+1)/2 \rceil = c + 1$ with $c = p/2$, so the $m = p + 1$ constraint is implied by $|S \cap P| \le c$ and is redundant in both orders. Hence $\mathcal{I}_{\sigma^-} = \mathcal{I}_{\sigma^+} = M_0$ and $B^- = B^+$.

Step 2: p odd. Then $\lceil p/2 \rceil = \lceil (p+1)/2 \rceil = (p+1)/2 = c$, and the $m = p+1$ constraint forbids exactly "$|S \cap P| = c$ and (first swapped item) $\in S$." Thus $\mathcal{I}_{\sigma^-}$ is M_0 with the capacity of $L := P \cup \{i\}$ lowered one unit (from its M_0-effective $c+1$ to c), and $\mathcal{I}_{\sigma^+}$ likewise lowers $P \cup \{j\}$. Let B_0 be the maximum-A basis of M_0.

(2a) $B^+ = B_0$. It suffices that $B_0 \in \mathcal{I}_{\sigma^+}$, i.e. B_0 never has both $|B_0 \cap P| = c$ and $j \in B_0$. Suppose $|B_0 \cap P| = c$. The $m = p+2$ constraint $|B_0 \cap (P \cup \{i, j\})| \le c + 1$ forces B_0 to contain at most one of i, j. If it contained j but not i, the swap $j \to i$ keeps every M_0 constraint (symmetry) and strictly raises A ($A_i > A_j$), contradicting maximality. So B_0 never holds j-without-i under P-saturation; hence $B_0 \in \mathcal{I}_{\sigma^+} \subseteq M_0$, and being M_0-optimal it is $\mathcal{I}_{\sigma^+}$-optimal: $B^+ = B_0$.

(2b) single swap. The order σ^- imposes on M_0 the single side constraint $|S \cap L| \le c$ with $L = P \cup \{i\}$, so B^- is the maximum-A basis of M_0 subject to $|S \cap L| \le c$. The retained $m = p+2$ constraint gives $|B_0 \cap (P \cup \{i,j\})| \le c+1$, hence $|B_0 \cap L| \le c+1$: the hypothesis of Lemma S3.5 holds with $r = c$. Therefore $B^- = B_0$, or $B^- = B_0 \setminus \{z\} \cup \{w\}$ with $z \in B_0 \cap L \subseteq P \cup \{i\}$, $w \in X \setminus L = \{j\} \cup T$, and $A_z > A_w$. Combining with (2a), $B^+ = B^-$, or $B^+ = B^- \setminus \{w\} \cup \{z\}$ with $z \in P \cup \{i\}$, $w \in \{j\} \cup T$, $A_z > A_w$.

(2c) ecological direction. $z \in P \cup \{i\}$ gives $Y_{\lambda_*}(z) \ge$ tied score; $w \in \{j\} \cup T$ gives $Y_{\lambda_*}(w) \le$ tied score; so $Y_{\lambda_*}(z) \ge Y_{\lambda_*}(w)$. With $A_z > A_w$ and $\lambda_* > 0$,

$$E_z = Y_{\lambda_*}(z)\, A_z^{\lambda_*} \ge Y_{\lambda_*}(w)\, A_z^{\lambda_*} > Y_{\lambda_*}(w)\, A_w^{\lambda_*} = E_w.$$

Therefore $\sum_{x \in B^+} E_x - \sum_{x \in B^-} E_x = E_z - E_w > 0$ (or $= 0$ if $B^+ = B^-$), giving $G^{\mathrm{BR}}(\sigma^+) \le G^{\mathrm{BR}}(\sigma^-)$. □

Proof of the theorem

We now prove Theorem 3.2 (stated in the main text, Section 6).

Proof of Theorem 3.2. By Lemma S3.3, $G_N^{\mathrm{BR}}(\lambda) = \sum_X E - \sum_{B_{\sigma(\lambda)}} E$ depends on λ only through the committed order $\sigma(\lambda)$. As in Theorem 3.1, there are at most $\binom{N}{2}$ interior crossings $0 < \lambda_1 < \cdots < \lambda_M < 1$; between consecutive crossings $\sigma(\lambda)$ is constant, so G_N^{BR} is constant. At a generic crossing only one pair $\{i, j\}$ ties. Every other plot has score strictly above or strictly below their common score at the crossing, and by continuity remains on the same side in a small neighbourhood. Hence the local order change is exactly the adjacent transposition $P\,i\,j\,T \mapsto P\,j\,i\,T$. Simultaneous pair crossings and higher-order ties would create block permutations, but these are precisely the non-generic events excluded by condition (iii) and are null events under the absolute-continuity condition above. Lemma S3.6 gives a weak decrease across this transposition. At the crossing point itself the fixed tie-breaker τ selects one of the two one-sided orders; since Lemma S3.6 orders the two one-sided values, either choice preserves weak monotonicity. Hence G_N^{BR} is weakly decreasing on $[0, 1]$. $\square$

Remark S3.7 (Verification audit)*.* The reduction and the exchange step were subjected to a hostile computational audit, reported in full in the replication archive; we summarise it here. Comparing the dynamic backward-induction solver against the prefix-matroid greedy basis on 119,000 fixed-order cases produced zero mismatches in realised set, Farmer payoff, Green payoff, or greedy-versus-enumeration basis. Across 344,893 adjacent-swap crossings there were zero failures of the exchange lemma, zero structural failures, and zero ecological-monotonicity failures. The parity split matched the proof exactly: even-prefix crossings were always basis-identical (Step 1), while odd-prefix crossings sometimes changed basis but always realised the single predicted exchange with the ecological inequality of Step 2 (Lemma S3.6). The audit is not used in the proof; its role is to stress-test the reduction and debug the formal argument, not to stand in for the theorem.

Remark S3.8 (Honest boundaries)*.* Three boundaries are worth stating explicitly. *(i) Claims World only.* The reduction relies on equal alternating claims and full allocation, which make the realisable family a matroid; the finite Budget-World counterexamples (Proposition S1.1) show this structure is destroyed once budgets bind, so the matroid route does not extend to Budget World. *(ii) Follower, not joint SPE.* The result concerns Farmer's subgame-perfect best response to a *committed* deterministic Green rule on the interval $[0, 1]$; the rule includes the exogenous tie-breaker τ at score ties. Green is not choosing among tied plots strategically or mixing at indifference points inside this result. The theorem is fully consistent with Hot Spot ($\lambda = -1$) outperforming value-first as a commitment rule (Appendix S2.4): that is a statement about $\lambda < 0$ and about rule commitment, not about the $[0, 1]$ tilt. *(iii) Genericity.* Condition (iv), no equal A-sums across Farmer bases, is what makes "the" SPE Farmer set well-defined;

without it Green's realised payoff can depend on Farmer's tie-breaking among equal-A bases.

Part II. Supplementary simulation results and robustness exercises

S4 Simulation Implementation Details

Each team $T \in \{T_f, T_g\}$ receives an integer claim endowment C_T, with $C_f + C_g = N$. The baseline is equal power, $C_f = C_g = N/2$, played over $R = 50$ rounds, with Farmers moving first. More generally, team T's claim schedule is an integer vector $(C_T^1, \dots, C_T^R)$ satisfying $\sum_{r=1}^{R} C_T^r = C_T$. When C_T/R is non-integer, the simulation assigns $\lfloor C_T/R \rfloor$ claims per round and distributes the remaining claims one per round starting from the earliest rounds, as in the unequal-power exercise of Appendix S8.

Let S_T^r be the set of plots claimed by team T in round r, and let

$$S_T^{\le r} = \bigcup_{k=1}^{r} S_T^k, \qquad U^{r,0} = P \setminus (S_f^{\le r-1} \cup S_g^{\le r-1})$$

be the cumulative claimed set and the set available at the start of round r. Within a round, teams claim sequentially without replacement, Farmers first and then Greens. We write $s_{f,\ell}^r$ and $s_{g,\ell}^r$ for the ℓth Farmer and Green claims and $U_f^{r,\ell-1}$, $U_g^{r,\ell-1}$ for the available sets just before those claims. The same priority rule is applied until the round quota is exhausted; all deterministic rules use the same fixed tie-breaker.

For the Rival-Aware heuristic, let $\tau = C_g/N$ and let $\mathcal{R}_\tau$ be the top $\lceil \tau N \rceil$ plots by environmental value. Farmer first claims the highest-a available plot inside this risky set, then reverts to descending a outside it:

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

S5 Claims World Baseline Conservation Achieved

Table S5.1 reports the exact final conservation scores for the three primary Green strategies against the two Farmer strategies at leakage levels of 0%, 50%, and 100%, decomposed into Pure Strategy Effect (PSE) and Displacement–Leakage Effect (DLE).

Table S5.1: Conservation Scores Under Different Strategies and Leakage Levels: The *Pure Strategy* and *Displacement-Leakage* Effects

Farmer	**Green**	**Leakage (%)**	**PSE**	**DLE**	**Final**
Rival-Blind	Max Environment	0	337	17	354
Rival-Blind	Max Environment	50	333	6	339
Rival-Blind	Max Environment	100	331	0	331
Rival-Blind	Block Farmers	0	252	126	378
Rival-Blind	Block Farmers	50	252	60	312
Rival-Blind	Block Farmers	100	252	0	252
Rival-Blind	Hot Spot	0	344	63	408
Rival-Blind	Hot Spot	50	354	29	383
Rival-Blind	Hot Spot	100	361	0	361
Rival-Aware	Max Environment	0	301	34	336
Rival-Aware	Max Environment	50	294	12	306
Rival-Aware	Max Environment	100	290	0	290
Rival-Aware	Block Farmers	0	142	7	148
Rival-Aware	Block Farmers	50	142	3	144
Rival-Aware	Block Farmers	100	142	0	142
Rival-Aware	Hot Spot	0	237	22	259
Rival-Aware	Hot Spot	50	241	11	252
Rival-Aware	Hot Spot	100	244	0	244

Notes. PSE = Pure Strategy Effect; DLE = Displacement–Leakage Effect; Final = PSE + DLE.

S6 Claims World: Supplementary Simulation Figures

This appendix collects the Claims-World simulation figures referenced from the main paper but not displayed there: static and dynamic Additionality; static and dynamic Welfare Loss; and the priority-tilt logic diagram. Static Conservation results appear in the main paper (Figure 3).

Final Additionality

Figure S6.1 reports final Additionality scores for all five Green strategies under both Farmer types. As discussed in Section 5.2.2, the qualitative ranking mirrors that of final Conservation, since the BAU benchmark is strategy-invariant across Green heuristics.

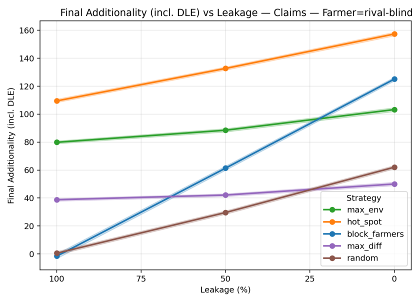

(a) Rival-Blind Farmers

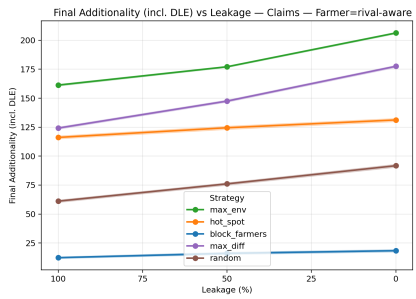

(b) Rival-Aware Farmers

Figure S6.1: Claims World: Final Additionality by Leakage

Additionality of Green Strategies: Dynamic Trajectories

Figure S6.2 reports the dynamic Additionality trajectories for all five Green strategies under both Farmer types; the static final-outcome figure above (Figure S6.1) and the detailed analytical discussion in Section 5.2.2 provide the static counterpart. As noted there, the trajectories differ from those for total conservation in one respect: additionality only increases when Greens capture a Farmer BAU plot, so end-of-game residuals (which are almost never BAU) generate no last-round jump — the trajectory of additionality is the trajectory of BAU-targeted captures alone.

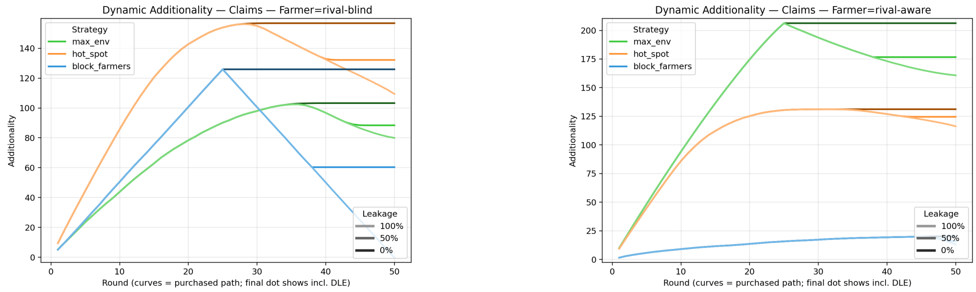

(a) Rival-Blind Farmers **(b)** Rival-Aware Farmers

Figure S6.2: Claims World: Dynamic Additionality Trajectories

Welfare Outcomes

Figure S6.3 reports the final Welfare-Loss statistic for all five Green strategies under both Farmer types; shaded bands indicate 95% bootstrap confidence intervals for the mean across Monte Carlo replications.

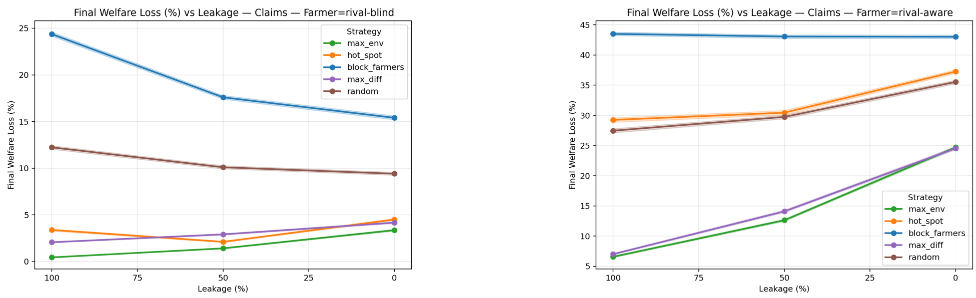

(a) Rival-Blind Farmers **(b)** Rival-Aware Farmers

Figure S6.3: Claims World: Final Welfare-Loss Statistic (%) by Leakage

Welfare Dynamic Trajectories

Figure S6.4 reports the dynamic Welfare Loss trajectories for the three primary Green strategies under both Farmer types. The mechanisms behind these trajectories are discussed in Section 5.2.3.

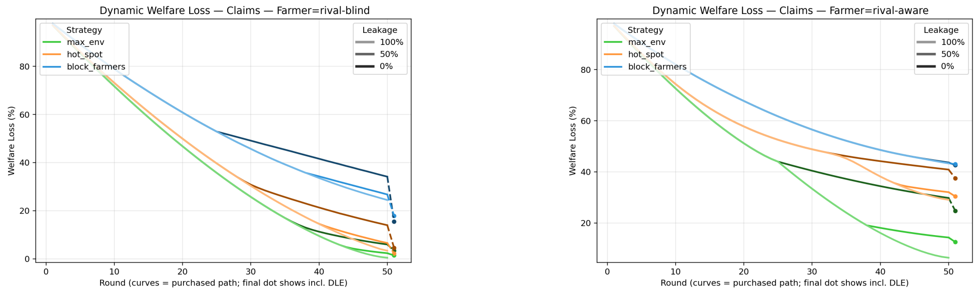

(a) Rival-Blind Farmers

(b) Rival-Aware Farmers

Figure S6.4: Claims World: Dynamic Welfare Loss Trajectories (%)

Logic of the Claims-World priority-tilt theorem

Figure S6.5 depicts the cumulative gap along the priority-tilt path $Y_\lambda = EA^{-\lambda}$, from value-first at $\lambda = 0$ to the ratio rule at $\lambda = 1$. The diagram visualises the layer-cake argument in the proof of Theorem 3.1 (Appendix A): each score crossing weakly lowers Green's value, and the endpoint inequality $G_N(0) \geq G_N(1)$ follows by finite induction over the at most $\binom{N}{2}$ crossings.

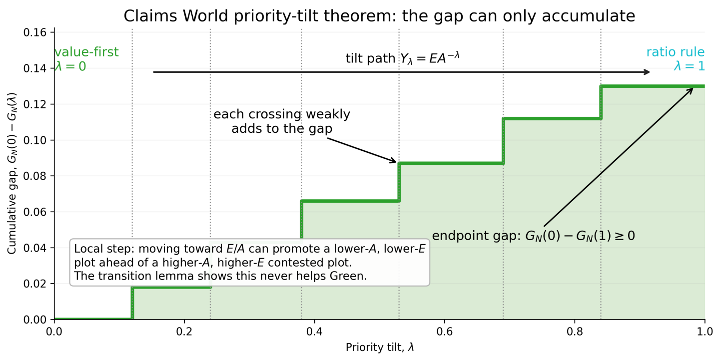

Figure S6.5: Logic of the Claims World priority-tilt theorem. The cumulative gap $G_N(0) - G_N(\lambda)$ along the priority-tilt path $Y_\lambda = EA^{-\lambda}$, from value-first at $\lambda = 0$ to the ratio rule at $\lambda = 1$. Vertical dotted lines mark score crossings; the endpoint gives $G_N(0) - G_N(1) \geq 0$.

S7 Budget World: Supplementary Simulation Figures

This appendix collects the Budget-World simulation figures referenced from the main paper but not displayed there: dynamic conservation trajectories under both Farmer types, and the full Welfare Loss results (final, purchased, and dynamic, under both Farmer types). Budget-World final and purchased Conservation results appear in the main paper (Figure 4).

Targeted-Denial Stress Test

In additional diagnostic runs, a costless policy-targeting adversary that deletes Green's next intended purchase largely eliminates the Max Environment advantage. The reversal therefore survives budgeted strategic denial but not free targeted denial. We treat this as a stress test rather than an economically meaningful Farmer strategy: a budgeted analogue would make the rival seek to harm Green rather than acquire land, outside the class of development-oriented behaviours studied here.

Budget World: Conservation Dynamic Trajectories

Figure S7.1 reports the dynamic Purchased Conservation trajectories under both Farmer types.

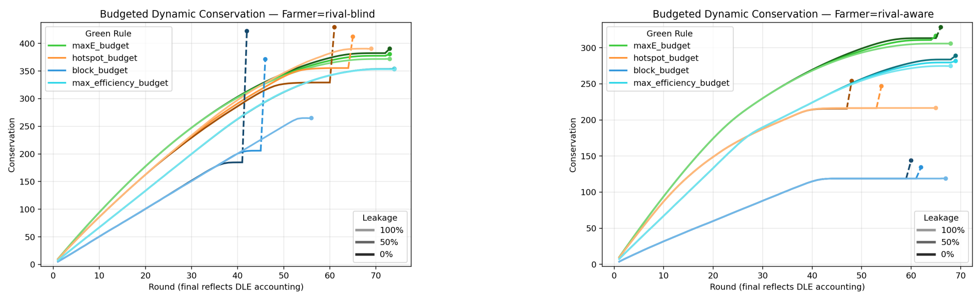

(a) Rival-Blind Farmers

(b) Rival-Aware Farmers

Figure S7.1: Budget World: Dynamic Conservation Trajectories

Budget World: Welfare Loss Outcomes

Figures S7.2 and S7.3 report final and purchased Welfare Loss together with dynamic trajectories under both Farmer types. As discussed in Section 5.3, the qualitative ranking parallels Claims World: Max Environment attains the lowest welfare losses, Hot Spot is close behind, and Block Farmers performs poorly. Welfare losses for the leading strategies typically increase as leakage falls, reflecting the same shifting-opportunity-set and residual-crediting mechanisms present in Claims World.

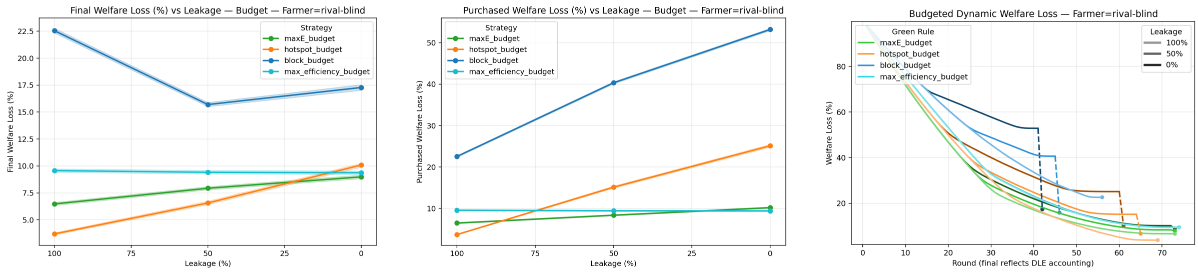

(a) Final (incl. DLE) Welfare Loss (%) **(b)** Purchased Welfare Loss (%) **(c)** Dynamic Trajectories

Figure S7.2: Budget World Green Strategies and Rival-Blind max-a Farmers: Social Welfare Loss (%)

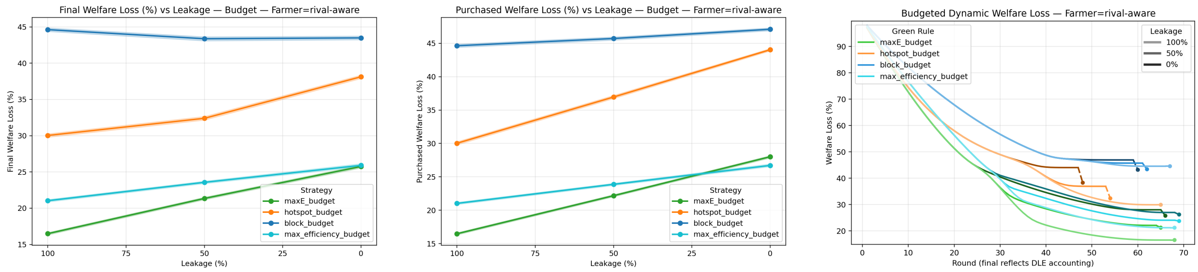

(a) Final (incl. DLE) Welfare Loss (%) **(b)** Purchased Welfare Loss (%) **(c)** Dynamic Trajectories

Figure S7.3: Budget World Green Strategies and Rival-Aware Farmers: Social Welfare Loss (%)

S8 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. The visible terminal movement is an accounting jump from residual-crediting (DLE) once Green has exhausted its claims, not a literal final-round burst of Farmer claims; the same end-of-game accounting affects all three outcome panels.

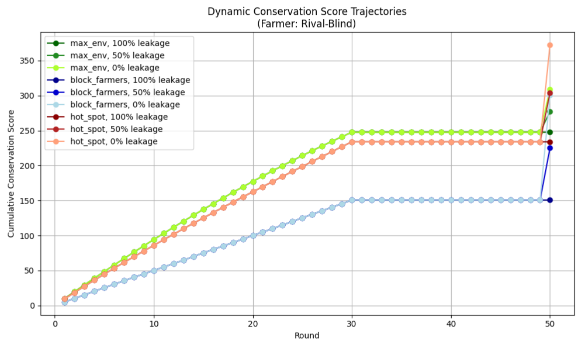

(a) Green Strategies and Rival-Blind Farmers

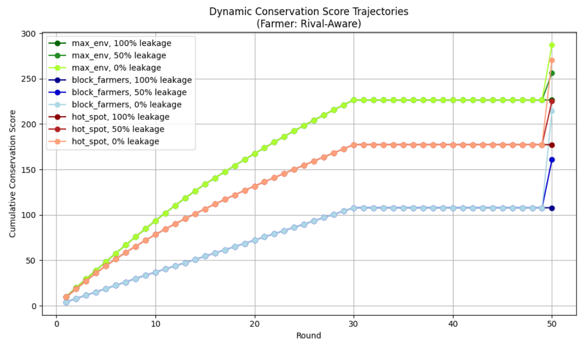

(b) Green Strategies and Rival-Aware Farmers

Figure S8.1: Conservation Scores for Political Allocation to Farmers = 70%

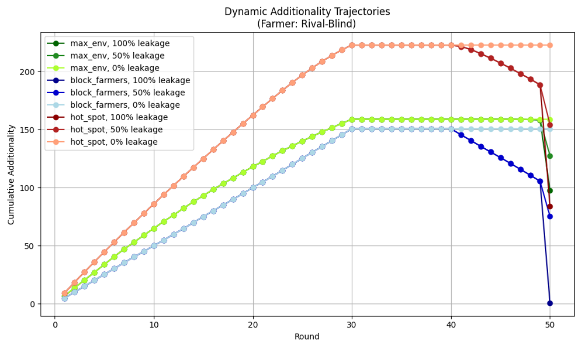

(a) Green Strategies and Rival-Blind Farmers

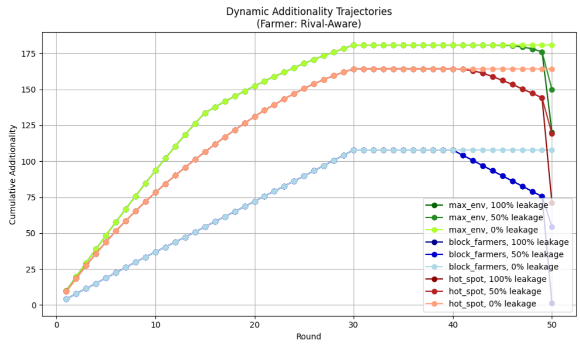

(b) Green Strategies and Rival-Aware Farmers

Figure S8.2: Additionality for Political Allocation to Farmers = 70%

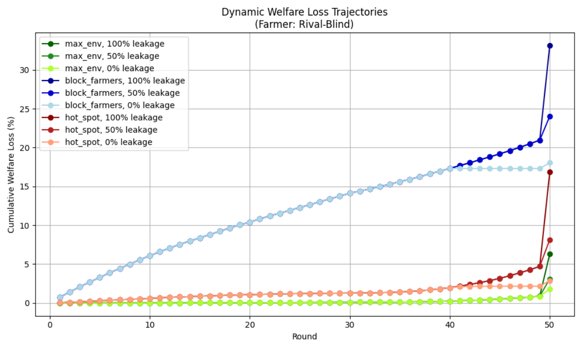

(a) Green Strategies and Rival-Blind Farmers

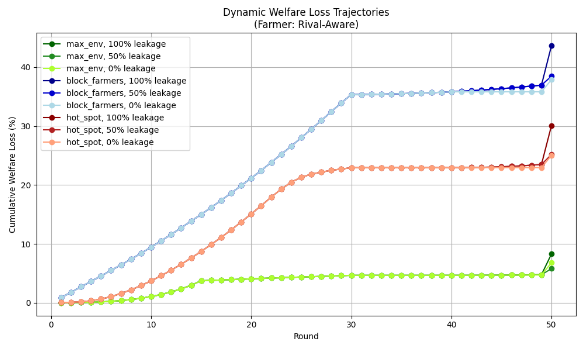

(b) Green Strategies and Rival-Aware Farmers

Figure S8.3: Social Welfare Loss(%) for Political Allocation to Farmers = 70%

S9 Spatial Externalities and Value-First Conservation

The baseline simulation grid is non-spatial: plots are exchangeable in their geographic structure. This appendix introduces a graph-based spatial spillover environment as a robustness lens and establishes one structural result — a connectivity dividend that aligns with ecological value and is orthogonal to agricultural value — that points value-first conservation in the same direction as the adversarial mechanism studied in the main text, through a separate channel.

Setup

Let $P = \{1, \ldots, N\}$ denote plots, and let $W = (W_{ij})$ be a non-negative symmetric adjacency matrix on P with $W_{ii} = 0$. The *effective* environmental value of plot i after some claims have been made is

$$e_i^{\mathrm{eff}}(S_g, S_f) \;=\; e_i \;+\; \eta_G \sum_{j \in S_g} W_{ij} \;-\; \eta_F \sum_{j \in S_f} W_{ij}, \tag{S9.1}$$

with $\eta_G, \eta_F \geq 0$. Green claims of plot j contribute a connectivity bonus $\eta_G W_{ij}$ to plot i's effective ecological value; Farmer claims contribute a fragmentation penalty. Direct ecological gains, spillover gains, and residual credits decompose cleanly:

$$\text{Final Conservation} \;=\; \underbrace{\sum_{i \in S_g} e_i}_{\text{PSE}^{\circ}} \;+\; \underbrace{\eta_G \sum_{i \in S_g} \sum_{j \in S_g} W_{ij} - \eta_F \sum_{i \in S_g} \sum_{j \in S_f} W_{ij}}_{\text{SEE (Spatial Externality Effect)}} \;+\; \underbrace{\text{DLE(Residual)}}_{\text{base-}e\text{ residual credit}},$$

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

The connectivity dividend

Define the set function on Green claims

$$F(S_g; S_f) = \sum_{i \in S_g} e_i + \eta_G \sum_{i \in S_g} \sum_{j \in S_g} W_{ij} - \eta_F \sum_{i \in S_g} \sum_{j \in S_f} W_{ij}, \tag{S9.2}$$

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

Proposition S9.1 (Connectivity dividend via supermodularity)**.** *If $W_{ij} \geq 0$, $W_{ij} = W_{ji}$, and $\eta_G \geq 0$, then $F(\cdot\,; S_f)$ is supermodular in S_g. The marginal gain from adding plot i to S_g is*

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

which is weakly increasing under set inclusion in S_g: if $S_g \subseteq S'_g$, then $\Delta_i(S_g; S_f) \leq \Delta_i(S'_g; S_f)$.

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 changes the Green–Green double sum by $2\sum_{j \in S_g} W_{ij}$, since both the new row and the new column contribute equally. As $W_{ij} \geq 0$, this cross-term is weakly increasing when S_g grows by set inclusion, which gives supermodularity. □

The proposition formalises a *connectivity dividend*: as a Green cluster grows, nearby plots become incrementally more attractive, independently of agricultural values a_i. The dividend is aligned with ecological value and orthogonal to agricultural value. It is therefore plausible that, among the fixed heuristics studied in this paper, value-first (E) tends to benefit relative to a-weighted rules such as Hot Spot (EA) and Block Farmers (A) as the connectivity dividend strengthens, because value-first selection rewards the same plots that connectivity rewards. We treat this as a heuristic comparative-static interpretation rather than a proved expected-payoff ranking theorem.

Relationship to the paper's main mechanism

The connectivity dividend operates through a separate channel from the adversarial knapsack reversal characterised in Theorem 3.1. The Claims-level theorem holds even with no spatial structure ($W = 0$); the supermodularity result here holds even with no rivalry (set $\eta_F = 0$ and consider a single agent). The two mechanisms point in the same direction — both favour value-first selection — but for distinct structural reasons. In landscapes where ecological and agricultural values are positively correlated and spatial spillovers are present, the two effects compound, and the empirically common conservation setting (heterogeneous spatial connectivity, moderate positive correlation) lies in the regime where they reinforce one another.

Part III. Bolivia supplement, replication, and methodological disclosures

S10 Supplement to Empirical Exercise on Bolivia

S10.1 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 illustrates 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. The protected-area layer in Figure S10.1(b) also distinguishes Indigenous jurisdictions; where these carry formal protection status in the source data they are included in the totals and shown separately in the legend. As illustrated in Figure S10.1(b), the dated protected-area series used here increases from zero in 1939 to 35.4 million hectares, 32% of the national territory[2]. Figure S10.1 reveals that PA expansion in Bolivia has proceeded at least three times faster than agropastoral expansion since 1985 (30 million hectares for conservation versus 8 million hectares for agropastoral expansion).

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

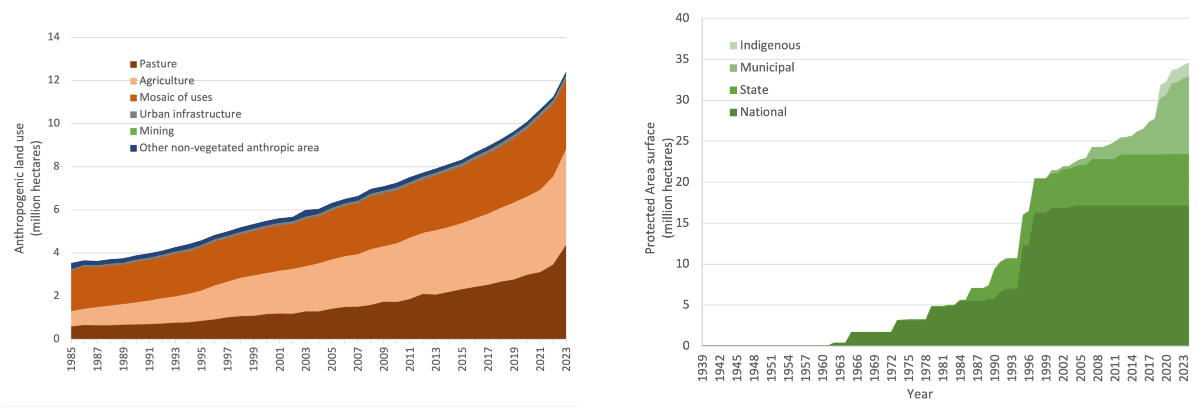

(a) Cumulative Anthropogenic Land Use **(b)** Cumulative Protected Areas

Figure S10.1: Cumulative land use in Bolivia: agriculture 1985–2023, protected areas 1939–2023

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, 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 maximising 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 S10.2 and as we illustrate below in figure S10.5, 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, recognising 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 S10.2 shows the dynamic expansion of agricultural land and Protected Areas from 1985 to 2024. Figures S10.3(a) and S10.3(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 S10.2: Supplementary Video: Dynamic expansion of Agricultural and Protected Area Land, 1985–2024
(Click the button to view the video).

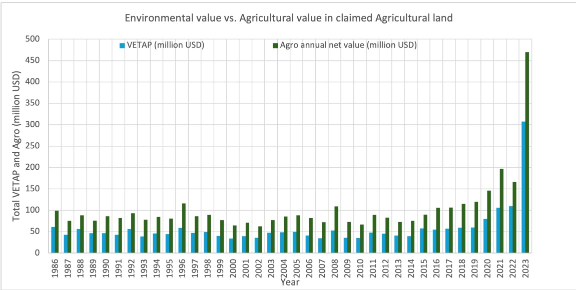

(a) New Agricultural Land, by year

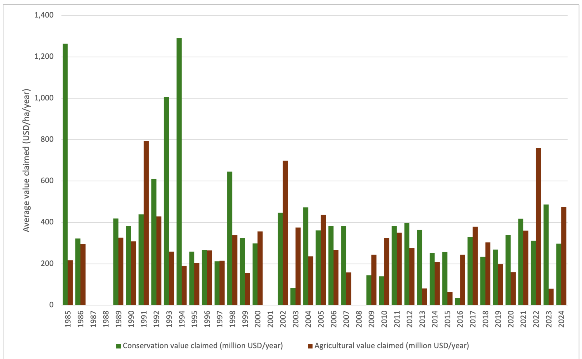

(b) New Protected Areas, by year

Figure S10.3: Potential Agricultural & Conservation Values of New Agricultural and Protected Land, (annual USD per ha)

S10.2 Conservation and agricultural expansion in Bolivia

In this section 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.

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 S10 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 S10.3 allows us to examine patterns in land allocation through the lens of our theoretical framework.

Specifically, Figure S10.5 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 S10.5 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.

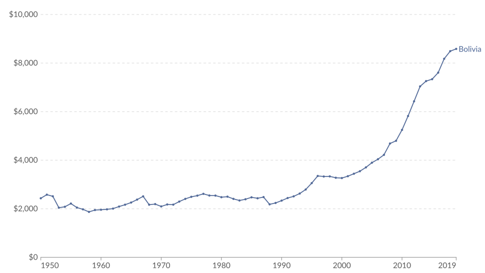

Figure S10.4: Bolivian GDP per capita, 1950–2019
(constant 2017 USD adjusted for cost of living)

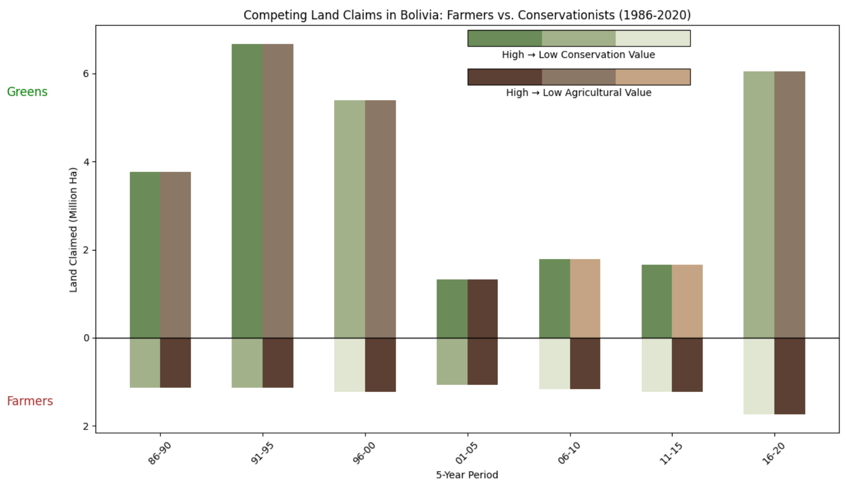

Figure S10.5: Agricultural & Conservation Potentials and Area of New Land allocated to Agriculture and Conservation, 5-year increments

Comparing the potential agricultural values (shades of brown) in Figure S10.5 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 maximisation. The patterns in conservation targeting appear more varied.

Figure S10.4 shows the evolution of GDP per capita over the same period. Comparing this with Figure S10.5 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 S10.2 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 S10) illustrates how enforcement challenges may create a negative correlation between agricultural pressure and realised conservation outcomes. The SPE diagnostics in Appendix S2 confirm that the value-first dominance pattern holds robustly under mild negative correlation, while the more extreme anti-concordant case shifts the optimal priority tilt toward ratio-greedy — a regime where the rival's deletion priorities and ecological values diverge so sharply that pre-emption becomes counterproductive.

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.

S10.3 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 S10.6(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 S10.6(a), following the methodology from Andersen et al. (2023). Agricultural potential varies greatly across Bolivia due to differences in topography, soil quality, and climate, with the most profitable regions being those with climates appropriate for the production of high value crops, like fruits and berries[3]. To generate the map, we use detailed pixel-level geographical data on slope, soil quality, precipitation, and temperature distribution to develop a high-resolution Production Cost Factor. Combining this with information on the most common crop (or livestock), average yields, and prices in each municipality to generate pixel-level estimates of agricultural potential[4].

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

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

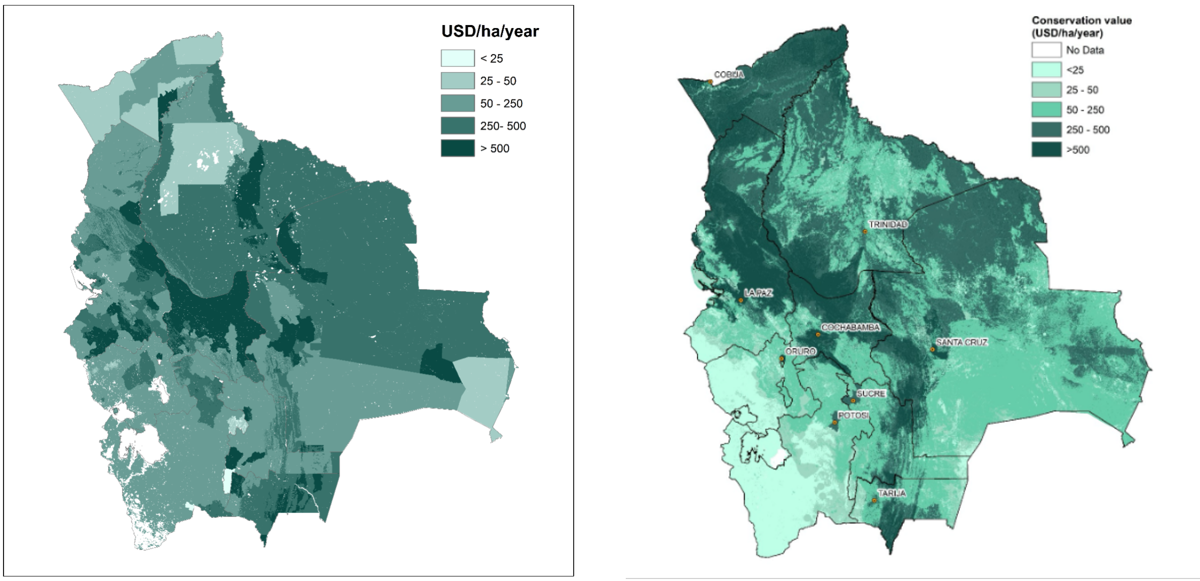

(a) Potential Agricultural Values **(b)** Potential Conservation Values

Figure S10.6: Potential Agricultural & Conservation values (annual USD per ha)

S10.4 Aggregating pixel values to planning units

Let Bolivia be discretised into raster pixels indexed by $i = 1, \ldots, I$. 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 $j = 1, \ldots, J$, and let $\mathcal{I}_j$ denote the set of pixels whose centres fall inside PU j.

Step 1: National pixel–level percentiles

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

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

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

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

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

Step 2: PU–level means and tail fractions

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

(S10.1)
$$\mu_j^{(e)} = \frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} e_i, \qquad \mu_j^{(a)} = \frac{1}{|\mathcal{I}_j|} \sum_{i \in \mathcal{I}_j} a_i,$$

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:

(S10.2)
$$f_j^{(e)}(q) = \frac{\#\left\{i \in \mathcal{I}_j : e_i \geq P_q^{(e)}\right\}}{|\mathcal{I}_j|}, \qquad q \in \{90, 95, 99\},$$

and analogously for agricultural value:

(S10.3)
$$f_j^{(a)}(q) = \frac{\#\left\{i \in \mathcal{I}_j : a_i \geq P_q^{(a)}\right\}}{|\mathcal{I}_j|}, \qquad q \in \{90, 95, 99\}.$$

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

Step 3: Tail–Averaged Percentile Scores (TAPS)

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

(S10.4)

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

and similarly for agricultural values:

(S10.5)

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

These scores emphasise PUs with a larger share of high-percentile pixels. The threshold labels $90, 95, 99$ are not themselves weights: because each $\mathrm{TAPS}_j(q)$ is converted separately into a percentile rank across PUs (Step 4), multiplying by these constants is order-preserving and does not affect the final ordering. Tail emphasis instead enters through the Step 5 component weights and the exponent on the 99th-percentile term.

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

(S10.6)
$$M_j^{(e)} \;=\; F_\mu^{(e)}(\mu_j^{(e)}),$$

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

$$T_j^{(e,90)},\ T_j^{(e,95)},\ T_j^{(e,99)} \tag{S10.7}$$

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

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

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

Step 5: Composite environmental and agricultural scores

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

$$S_j^{(e)} = 0.40\, M_j^{(e)} + 0.25\, T_j^{(e,90)} + 0.20\, T_j^{(e,95)} + 0.15\, \big(T_j^{(e,99)}\big)^{1.5}, \tag{S10.8}$$

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

We define the composite agricultural score analogously:

$$S_j^{(a)} = 0.40\, M_j^{(a)} + 0.25\, T_j^{(a,90)} + 0.20\, T_j^{(a,95)} + 0.15\, \big(T_j^{(a,99)}\big)^{1.5}. \tag{S10.9}$$

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.

S10.4.1 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 7(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 S10.4):

- **Environmental Value ($S_j^{(e)}$):** We utilise 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 normalised to a $[0, 1]$ scale based on their percentile ranks among all PUs, standardising the "cost" of a PU as a measure of its opportunity cost relative to the national average. Thus we move the analysis from a stylised grid to a board characterised by realistic heterogeneity: the underlying raster values and hotspot concentrations are highly skewed (the bounded, rank-normalised game scores themselves are not heavy-tailed in the cardinal sense), agricultural values follow complex soil and transport gradients, and the two values are positively correlated (with $\rho = 0.29$ on this index board, close to the $\rho = 0.3$ positive-correlation case analysed in Appendix S1.8; the percentile-rank transform raises the measured correlation, which is $\rho \approx 0.17$ on cardinal PU means). The ecological selection premium on this board is 22.7%, above the 12–15% threshold identified in Section 5.4, and rises to 74.3% when the same landscape is valued on cardinal PU means, so Bolivia lies in the region where the knapsack reversal operates on either scale (Appendix Table S10.1). The distribution of agricultural and environmental values is displayed in the heatmaps in Figure 7 (in the main paper).

S10.5 Value-scale robustness

The baseline board values each PU by composite percentile-rank indices $S_j^{(e)}, S_j^{(a)} \in [0,1]$ (Steps 1–5 above), so it is an index-board implementation of the two-value model rather than a cardinal price-board calibration. Because the ratio-greedy rule ranks on $S^{(e)}/S^{(a)}$, a quantity that is not invariant to a monotone rescaling of either coordinate, we check that the reversal is not an artefact of the percentile-rank normalisation by rerunning the Bolivia budget game on four alternative value scales: the raw cardinal PU means $\mu_j^{(e)}, \mu_j^{(a)}$, a positive min–max rescaling of those means, a log transform, and area-weighted raw totals (each PU's cardinal means scaled by its area in km^2).

Table S10.1 reports the results. The reversal survives on every scale. Raw cardinal means and the positive min–max rescaling strengthen it relative to the index board (Max Environment exceeds strict ratio-greedy Max Efficiency by 12.2–12.4% and 11.5–12.1% respectively in purchased conservation, against 4.5–7.6% on the index board), and the static ratio plan's overstatement rises correspondingly to 27.4–30.0%. Area-weighted totals strengthen it further still, though their larger magnitudes partly reflect the wide dispersion in basin area. The advantage tracks the ecological selection premium across the scales: the deliberately tail-compressing log transform drives the premium to 13.7%, inside the diagnostic threshold region, and attenuates the Max Environment advantage to 1.6–5.0%, but the reversal sign is preserved at every leakage rate. Finally, the exact 0–1 knapsack solution is virtually identical to the static ratio-greedy benchmark on every scale, with gains below 0.02%, so the static–dynamic gap reported throughout is an adversarial-procurement effect rather than approximation error in the passive benchmark.

Table S10.1: Bolivia value-scale robustness

Scale	$\rho(e, a)$	ESP (%)	ME over MX (%)	Static–dynamic MX gap (%)	ILP gain (%)
Rank-index board	0.29	22.7	4.5–7.6	19.1–21.0	0.02
Raw PU means	0.17	74.3	12.2–12.4	27.4–30.0	0.01
Positive min–max means	0.17	102.2	11.5–12.1	29.9–31.6	0.01
Log PU means	0.16	13.7	1.6–5.0	7.0–11.8	0.02
Area-weighted raw totals	0.75	416.6	49.3–54.4	44.2–44.9	0.01

Notes: Each row reruns the Bolivia budget game on the same 390 planning units with the indicated value scale. ME over MX is the percentage advantage of Max Environment over strict ratio-greedy Max Efficiency in purchased conservation under the rival-aware farmer, reported as the range across leakage rates $L \in \{1, 0.75, 0.5, 0.25, 0\}$. The static–dynamic MX gap is $100 \times (E^{\text{static ratio}} - E^{\text{dynamic MX}})/E^{\text{static ratio}}$, again for purchased conservation under the rival-aware farmer. ESP is computed on the same value scale as the simulation. ILP gain is the percentage improvement of the exact static knapsack solution over the static ratio-greedy solution.

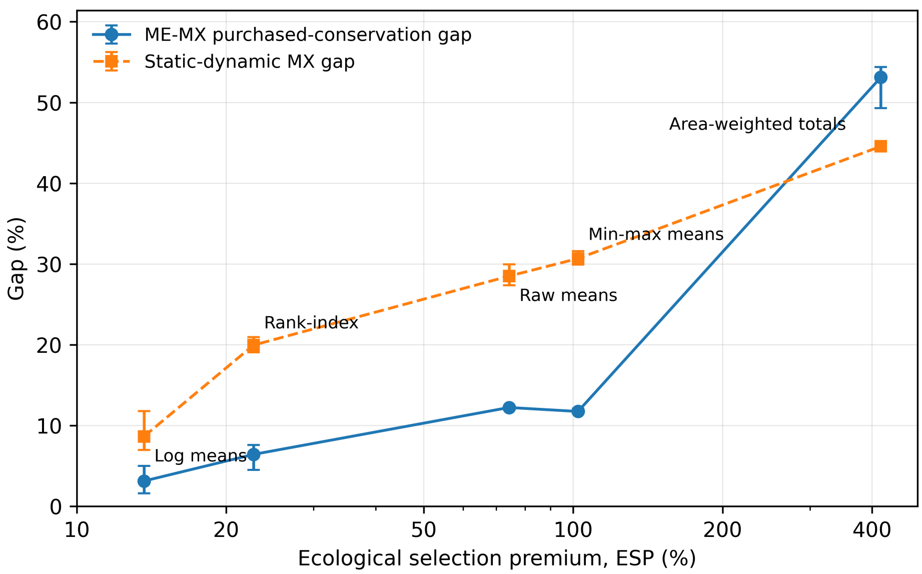

Figure S10.7: Ecological selection premium and dynamic gaps across Bolivia value scales. Each point reruns the same 390 planning units on the indicated value scale. The horizontal axis reports the ecological selection premium, computed on that same scale. Circles report the percentage advantage of Max Environment over strict ratio-greedy Max Efficiency in purchased conservation. Squares report the static–dynamic Max Efficiency gap, defined as the percentage by which the static ratio-greedy plan overstates realised dynamic Max Efficiency conservation. Markers are means across leakage rates $L \in \{1, 0.75, 0.5, 0.25, 0\}$ for the rival-aware farmer; vertical bars show the corresponding range. The horizontal axis is logarithmic because the area-weighted raw-total board has a much larger ESP than the equal-PU boards.

S10.6 Bolivia Case Study: Supplementary Simulation Figures and Maps

This appendix reports the full Bolivia-board results summarised in Section 7. We run both the Claims- and Budget-World contests on the fixed board of 390 Planning Units (PUs), under Rival-Blind and Rival-Aware Farmers and across the full leakage grid. Conditional on this single board, the results are deterministic. The figures below report static Conservation and Welfare-Loss outcomes, the corresponding dynamic trajectories and PSE/DLE decomposition, and maps of the spatial distribution of conserved PUs. Additionality is not reported separately because, with the strategy-invariant business-as-usual benchmark used here, its rankings reproduce those for Conservation.

Claims World: Static Conservation and Welfare Outcomes

Figures S10.8 and S10.9 report static final Conservation and Welfare-Loss outcomes on the Bolivia board for the three primary Claims-World Green strategies under both Farmer types. Against Rival-Blind Farmers, Hot Spot leads final Conservation at every leakage rate. Max Environment is next at high and intermediate leakage, while Block Farmers overtakes it at the two lowest leakage rates only because residual Conservation becomes large. Against Rival-Aware Farmers, Max Environment dominates throughout, Hot Spot is second, and Block Farmers performs poorly. The same mechanism as on the synthetic boards drives these rankings: high-e, high-a sites are either secured early by a value-led rule or lost to development.

The assignment-loss statistic supplies a distinct check on selection quality. Under both Farmer behaviours it is lowest under Max Environment and Hot Spot and much higher under Block Farmers. For the value-led rules, stricter leakage control generally raises the terminal loss even as it raises final Conservation. Thus preventing re-targeting can create residual Conservation without improving which side receives the plots it values most, reproducing the welfare-loss pattern from the large-board Monte Carlo simulations.

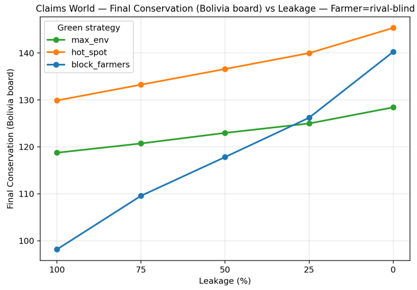

(a) Rival-Blind Farmers

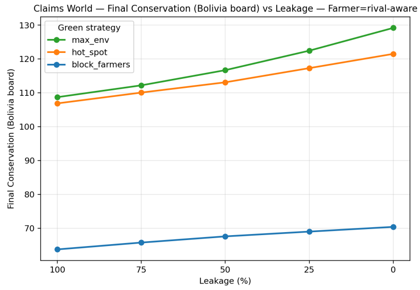

(b) Rival-Aware Farmers

Figure S10.8: Claims World on a Bolivian Board: Final Conservation Scores (incl. DLE)

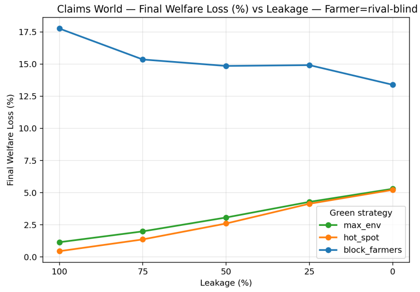

(a) Rival-Blind Farmers

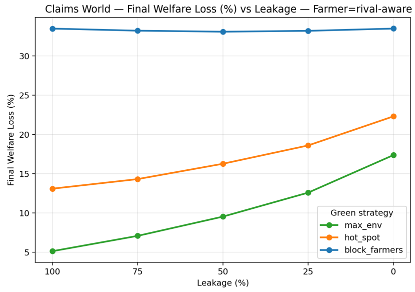

(b) Rival-Aware Farmers

Figure S10.9: Claims World on a Bolivian Board: Final Social Welfare Loss (%)

Budget World: Static Conservation

Figure S10.10 reports final Budget-World Conservation, including DLE, under both Farmer types. Against Rival-Blind Farmers, Hot Spot leads at most leakage rates, with Max Environment close behind. Block Farmers rises sharply as leakage falls and approaches the leaders at the lowest leakage rates, but the trajectories below show that this recovery reflects residual crediting rather than a stronger purchased portfolio. Against Rival-Aware Farmers, Max Environment clearly dominates, Max Efficiency remains below it at every leakage rate, and Block Farmers is weakest by a wide margin. The Bolivia board therefore reproduces the central selection result of the stylised Budget World: value-first retains its advantage over ratio-greedy when the opportunity set is generated by an active rival.

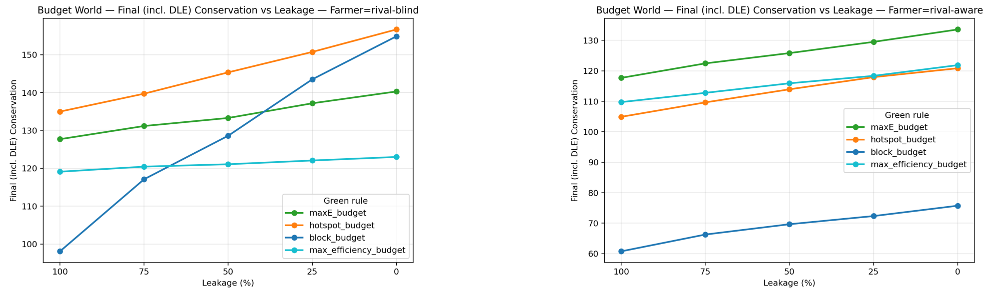

(a) Rival-Blind Farmers **(b)** Rival-Aware Farmers

Figure S10.10: Budget World on a Bolivian Board: Final Conservation Scores (incl. DLE)

Dynamic Trajectories

Figures S10.11–S10.13 make the underlying paths explicit. The solid trajectories record the ecological value Green directly secures—the PSE in Claims World and Purchased Conservation in Budget World—while the final dashed segments and dots add DLE. In Claims World against Rival-Blind Farmers, Max Environment initially accumulates ecological value faster, but Hot Spot catches and overtakes it by securing contested high-e, high-a plots. Against Rival-Aware Farmers, Max Environment leads throughout because departing from value-first exposes the highest-e plots to pre-emption. Block Farmers has the weakest direct-selection path; its low-leakage improvement in final Conservation comes from the DLE jump at the endpoint.

Budget World sharpens the same distinction. Against Rival-Blind Farmers, Hot Spot's Purchased Conservation path lies only moderately above Max Environment, and much of its low-leakage final edge is residual. Block Farmers' purchased portfolio remains weak; its apparent recovery is DLE. Against Rival-Aware Farmers, Max Environment leads on both Purchased and final Conservation. The nearly parallel Max Environment and Max Efficiency paths preserve the value-first advantage at every leakage rate, reproducing the selection effect analysed in Appendix S1.8. The Welfare-Loss trajectories show the complementary paradox: stricter leakage control can raise terminal assignment loss even as it increases Conservation.

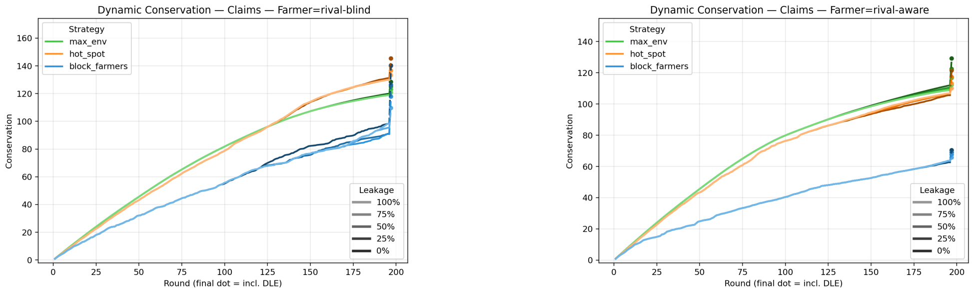

(a) Rival-Blind Farmers **(b)** Rival-Aware Farmers

Figure S10.11: Claims World on a Bolivian Board: Dynamic Conservation Trajectories

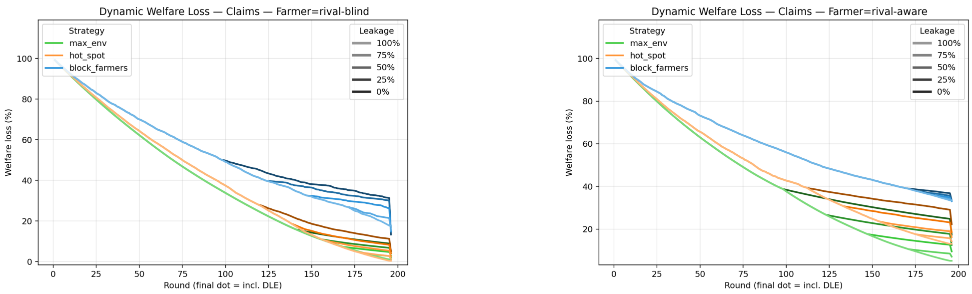

(a) Rival-Blind Farmers **(b)** Rival-Aware Farmers

Figure S10.12: Claims World on a Bolivian Board: Dynamic Welfare-Loss Trajectories (%)

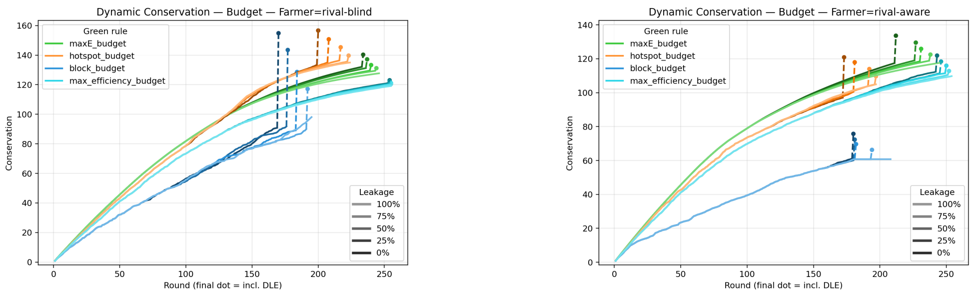

(a) Rival-Blind Farmers **(b)** Rival-Aware Farmers

Figure S10.13: Budget World on a Bolivian Board: Dynamic Conservation Trajectories

Spatial Distribution of Conserved PUs by Strategy

Figures S10.14, S10.15, and S10.16 display the spatial distribution of Planning Units conserved by each Green strategy on the Bolivian board (Budget World, leakage = 100%). The maps illustrate *where* each strategy locates its protected portfolio across Bolivia's geography and are reported for geographic comparison and transparency. Detailed interpretation is left to readers familiar with Bolivian biogeography.

(a) against Rival-Blind Farmers

(b) against Rival-Aware Farmers

Figure S10.14: Budget World Max Environment Strategy Conserved PUs (leakage=100%)

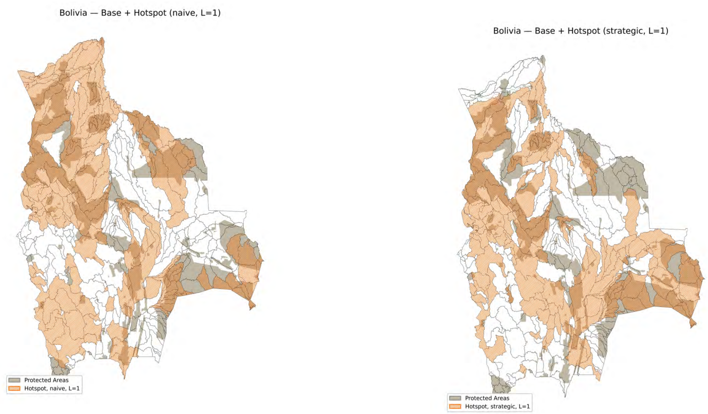

(a) against Rival-Blind Farmers

(b) against Rival-Aware Farmers

Figure S10.15: Budget World Hot Spot Strategy Conserved PUs (leakage=100%)

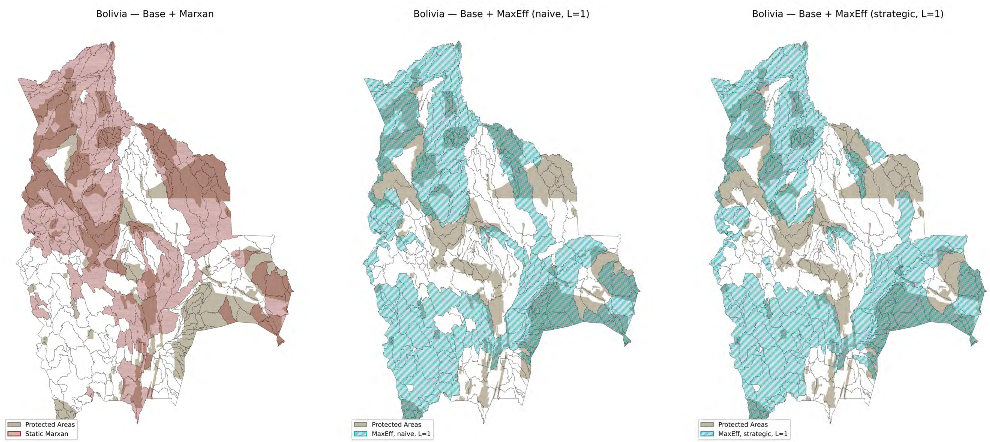

(a) Static Knapsack

(b) Dynamic Max Efficiency Rival-Blind

(c) Dynamic Max Efficiency Rival-Aware

Figure S10.16: Static Knapsack Benchmark and Dynamic Max Efficiency Conserved PUs (Budget World, leakage=100%)

S11 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 S11.1: 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.

S12 AI-Assisted Methodology

This paper was developed with substantial assistance from large language models in differentiated roles. Anthropic's Claude Opus contributed to drafting, mathematical development, proof checking, and structural editing. OpenAI's ChatGPT Pro contributed to formal proof development, computational verification, and adversarial auditing. Google's Gemini Pro contributed hostile auditing of formal claims. A companion paper documenting the human–AI collaboration methodology is available as Weinhold (2026b). The authors take sole responsibility for all claims, all errors, and all editorial decisions made in the paper.

References

Genz, A. (2004). Numerical computation of rectangular bivariate and trivariate normal and t probabilities. *Statistics and Computing* 14(3), 251–260. https://doi.org/10.1023/B:STCO.0000035304.20635.31.

Plackett, R. L. (1954). A reduction formula for normal multivariate integrals. *Biometrika* 41(3–4), 351–360. https://doi.org/10.1093/biomet/41.3-4.351.

Pollard, D. (1984). *Convergence of Stochastic Processes*. Springer, New York.

van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge University Press.

Shorack, G. R. and Wellner, J. A. (1986). *Empirical Processes with Applications to Statistics*. Wiley, New York.

Kurtz, T. G. (1970). Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability* 7(1), 49–58.

Ethier, S. N. and Kurtz, T. G. (1986). *Markov Processes: Characterisation and Convergence*. Wiley, New York.

Wormald, N. C. (1999). The differential equation method for random graph processes and greedy algorithms. In *Lectures on Approximation and Randomised Algorithms* (M. Karoński and H.-J. Rucinski, eds.), pp. 73–155. PWN, Warsaw.

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. https://sdsnbolivia.org/working-paper-n-5-2023-map-of-agricultural-potential-in-bolivia

Andersen, L. E., Argondoña, 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.

MapBiomas Bolivia Project (2024). MapBiomas Bolivia Collection 2.0: Land use/land cover and transitions 1985 to 2024. https://bolivia.mapbiomas.org/en/. Accessed March 2025.

SDSN Bolivia (2025). Áreas Protegidas de Bolivia, 2025. SDSN Bolivia – UPB, Datos Espaciales. https://sdsnbolivia.org/datos-espaciales/.

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. https://ourworldindata.org/grapher/gdp-per-capita-penn-world-table. Retrieved March 13, 2025.

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

Weinhold, D. (2026a). 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

Weinhold, D. (2026b). How a non-theorist and two AIs proved a theorem: anatomy of a human–AI collaboration in mathematical economics. Available at https://ssrn.com/abstract=6591059

R. A. Brualdi, Comments on bases in dependence structures, *Bull. Austral. Math. Soc.* 1 (1969) 161–167.