# A Game-Theoretic and Dynamical-Systems Framework for Anti-Poaching Resource Allocation: A Case Study of Etosha National Park

Chan Ka Hin[1], Ao Long Nam[1], Loi Weng Kin[1], Tse Long Tin[1], Cheng Sok Kin[1,†]

[1]Colégio Diocesano de São José (5ª), Macau

[†]Supervising instructor; correspondence: chengsokkin@cdsj5.edu.com

**Abstract**

Wildlife poaching threatens biodiversity across sub-Saharan Africa, and is especially acute for critically endangered species such as the black rhinoceros (*Diceros bicornis*). Etosha National Park, Namibia (22,935 km$^2$), is patrolled by approximately 295 anti-poaching rangers, posing two interlinked operational questions: *where* should a limited workforce be placed to maximise conservation return, and *how many* rangers constitute a dynamically viable minimum below which the prey population tips toward long-run decline?

We present a three-layer mathematical pipeline that addresses both questions jointly, with three analytical extensions beyond sequential layer application. Layer 1 generates a Wildlife Protection Potential (WPP) field combining nonlinear Multi-Attribute Utility Theory (MAUT) utilities for endangered (exponential) and abundant (logarithmic) species, a waterhole attraction term, and a non-homogeneous Poisson process (NHPP) threat intensity field. Layer 2 solves a Stackelberg-structured patrol allocation as a nonlinear constrained optimisation (SLSQP) over the WPP field, incorporating a season-adaptive terrain-friction matrix. Layer 3 analyses a three-population prey–predator–poacher ODE system; Jacobian eigenvalue analysis identifies the critical ranger threshold $\mathcal{T}^*$.

Three extensions advance beyond the base pipeline. First, a *self-consistent feedback loop* iterates the three layers to a fixed-point equilibrium in which the deployment is optimal given the threat landscape and the threat landscape is consistent with the population dynamics induced by that deployment. Second, a *multi-objective Pareto analysis* maps the trade-off between capture efficiency and patrol equity, identifying the knee point of maximum marginal return. Third, real *Diceros bicornis* occurrence records from GBIF ($n = 103$ within the Etosha bounding box) are integrated as a spatial prior, anchoring the WPP field to observed animal distributions.

Applied to Etosha, the framework yields a patrol allocation (Gini $\approx$ 0.77) anchored on waterhole arcs, a stability threshold of approximately 240–260 rangers, and a feedback-loop objective gain of $\approx$ 33% over a single open-loop pass. Sensitivity analysis confirms that UAV and sensor uptime dominate over headcount adjustments.

**Keywords:** anti-poaching; Stackelberg security game; multi-objective optimisation; Lotka–Volterra; terrain friction; GBIF occurrence data; self-consistent equilibrium; Etosha National Park.

# 1 Introduction

## 1.1 Background and context

Wildlife poaching is a persistent threat to global biodiversity, and especially acute in sub-Saharan Africa, where critically endangered species such as the black rhinoceros (*Diceros bicornis*) face intense illegal harvest pressure.[1,2] Etosha National Park in Namibia illustrates the operational challenge: roughly 22 935 $km^2$ of heterogeneous savannah and salt pan is patrolled by approximately 295 anti-poaching rangers—fewer than 0.02 per $km^2$. Such ratios are common across large African reserves and make uniform patrol deployment wasteful. A concentrated subset of sites—permanent waterholes, calving grounds, known rhino home ranges— simultaneously concentrates ecological value and poaching pressure.[3]

## 1.2 Problem statement

Two interlocking questions arise that are often treated separately. *Where* should a limited workforce be placed to maximise conservation return, given uneven terrain and adaptive adversaries? And *how many* rangers constitute a dynamically viable minimum? We argue that answering either question without the other risks producing a deployment that is spatially efficient but dynamically unsustainable, or a threshold estimate with no spatial content. This paper proposes a three-layer pipeline that addresses both from the same inputs, with three novel extensions: a self-consistent feedback mechanism connecting Layers 1–3, a multi-objective Pareto analysis, and real-data spatial anchoring via GBIF occurrence records.

## 1.3 Related work

**Game-theoretic anti-poaching.** PAWS[4] established Stackelberg Security Games (SSGs) for patrol allocation; CAPTURE[5] extended this with predictive Bayesian modelling, validated by a five-month field trial in Uganda.[6] A core challenge is that sparse incident records are biased toward already-patrolled areas. We address this through a spatial prior derived from real occurrence data, in keeping with the recommendation of Esfahani and Kuhn[7] to anchor models to the best available data.

**Ecological valuation and spatial planning.** Multi-criteria prioritisation is motivated by Billionnet;[8] Moffett et al.[9] apply MAUT to north-central Namibia directly. Redfern et al.[3] established that surface-water availability dominates large-herbivore distribution in semi-arid savannah, making waterholes natural anchors for any priority map. Terrain connectivity is reviewed in Moilanen et al.[10]

**Technology and real-data anchoring.** Koh and Wich[11] proposed UAVs as coverage enhancers; Wall et al.[12] demonstrated real-time sensor networks. We leverage GBIF occurrence records[13] as an observational anchor, complementing sensor-network data with the published biodiversity database.

**Dynamical models.** Milner-Gulland and Leader-Williams[1] and Bulte and van Kooten[2] developed economic poaching models. Lotka–Volterra prey–poacher formulations have been used to analyse stability conditions.[14] We extend these to a three-population system and couple the dynamical analysis to the spatial allocation via a self-consistent iteration, a link absent from prior work.

**Contribution.** We integrate spatial Stackelberg game theory, nonlinear MAUT, three-population ODE stability analysis, real-data spatial priors, multi-objective Pareto analysis, and a self-consistent feedback mechanism in a single transparent pipeline. Every parameter has an explicit value; every modelling choice can be scrutinised and varied.

Section 2 presents the methods. Section 3 reports the Etosha application. Section 4 discusses limitations and implications.

# 2 Methods

## 2.1 Model assumptions

**A1. Boundedly rational poachers.** Poachers prefer cells with higher ecological value and lower detection probability, but hold imperfect information about patrol schedules. This justifies the Stackelberg leader–follower structure without requiring omniscient adversaries.

**A2. Terrain friction dominates travel cost.** Ranger logistical cost is proportional to a friction coefficient derived from terrain type, not Euclidean distance. The Etosha Pan becomes impassable mud in the wet season; this is captured by a seasonal friction multiplier.

**A3. A three-population ODE suffices for the stability question.** We model the ecosystem as prey (rhino) $N$, natural predator (lion/leopard) $P$, and poacher $Z$. The average patrol density $\bar{\tau} = \mathcal{T}/A$ enters the ODE as a fixed parameter, appropriate for asking whether a given headcount sustains a stable equilibrium.

**A4. Technology modifies detection efficiency.** UAVs and sensor networks are modelled as altering the detection efficiency $\lambda_{\text{det}}$: full technology gives $\lambda_{\text{det}} = 0.25$; degraded technology gives $\lambda_{\text{det}} = 0.08$.

## 2.2 Notation

**Table 2.1:** *Principal notation.*

| Symbol | Description | Type / Range |
|---|---|---|
| *Operational* | | |
| $\mathcal{T}$ | Total ranger workforce | Positive integer |
| $\mathcal{T}^*$ | Stability threshold | Output |
| $A$ | Park area (22,935 $\text{km}^2$) | Fixed |
| $\bar{\tau} = \mathcal{T}/A$ | Average patrol density | Rangers per $\text{km}^2$ |
| $\tau_i$ | Rangers allocated to cell $i$ | $\geq 0$ |
| $\lambda_{\text{det}}$ | Detection efficiency | $> 0$ |
| $\mu_i$ | Terrain-friction coefficient | $\geq 1$ |
| *WPP field (Layer 1)* | | |
| $U_R(N)$ | Endangered species utility | $[0, 1]$ |
| $U_E(N)$ | Abundant species utility | $[0, 1]$ |
| $w_R$, $w_E$ | Conservation priority weights | 0.65, 0.35 |
| $\lambda(s, t)$ | NHPP threat intensity | $\geq 0$ |
| $\boldsymbol{\beta}$ | NHPP covariate weights | $\mathbb{R}^4$ |
| $W_i$ | WPP score for cell $i$ | $\geq 0$ |
| $\mathcal{R}_i(\tau_i)$ | Interception probability | $[0, 1]$ |
| *Population dynamics (Layer 3)* | | |
| $N(t)$, $P(t)$, $Z(t)$ | Prey, predator, poacher densities | $\geq 0$ |
| $r$, $K$ | Prey growth rate; carrying capacity | Parameters |
| $\alpha_p$, $\alpha_z$ | Predation and poaching rates | Parameters |
| $\eta$, $\delta$ | Predator conversion, mortality | Parameters |
| $\rho$, $m$, $\sigma$ | Poacher recruitment, exit, suppression | Parameters |
| $J$ | Jacobian at interior equilibrium | $3 \times 3$ real |

## 2.3 Layer 1: WPP field via nonlinear MAUT and NHPP

### Utility functions

Two species groups carry different conservation urgency. For the critically endangered black rhinoceros, protection value rises steeply once the population exceeds a critical viability floor $N_{\text{crit}}$; for abundant species (elephant), value exhibits diminishing marginal returns.

The *endangered* utility:

$$U_R(N) = 1 - \exp\Big(-\alpha_U \frac{\max(N - N_{\text{crit}},\, 0)}{N_{0,R}}\Big), \tag{1}$$

with $\alpha_U = 0.08$ and $N_{0,R} = 300$. This convex shape near $N_{\text{crit}}$ encodes the all-or-nothing importance of remaining above the viability floor.

The *abundant* utility:

$$U_E(N) = \frac{\ln(1 + \kappa N)}{\ln(1 + \kappa N_{0,E})}, \tag{2}$$

with $\kappa = 0.02$ and $N_{0,E} = 2500$. The logarithmic shape captures diminishing returns with increasing population size.

**Waterhole attraction and NHPP threat**

Waterholes concentrate both wildlife and poaching activity.[3] We encode this with a spatial attraction term over all $H = 86$ permanent waterholes:

$$A_i = \sum_{h=1}^{H} 10 \, \exp\big(-0.5\, d_{ih}\big), \tag{3}$$

where $d_{ih}$ is the Euclidean distance (grid cells) from cell $i$ to waterhole $h$.

The threat landscape is modelled as a non-homogeneous Poisson process (NHPP) with intensity:

$$\lambda(s,t) = \lambda_0(t) \, \exp\big(\boldsymbol{\beta}^\top \mathbf{X}(s,t)\big), \tag{4}$$

where $\lambda_0(t) \in \{0.55, 0.35\}$ (dry/wet season) and the covariate vector $\mathbf{X} = (d_{\text{water}}, d_{\text{edge}}, d_{\text{road}}, \mu)$ captures proximity to water, park boundary, road, and terrain friction, with $\boldsymbol{\beta} = (-0.45, -0.25, -0.15, -0.18)^\top$.

**Composite WPP field**

The Wildlife Protection Potential for cell $i$ is:

$$W_i = \big(w_R \, U_R(N) + w_E \, U_E(N)\big) \cdot A_i \cdot \big(1 + \lambda(s_i, t)\big), \tag{5}$$

with $w_R = 0.65$ and $w_E = 0.35$, reflecting the relative management priority of the black rhinoceros (critically endangered) versus elephant (abundant in Etosha).[1]

**GBIF spatial prior**

Publicly available *Diceros bicornis* occurrence records from GBIF[13] anchor the WPP field to observed distributions:

$$W_i^{\text{prior}} = W_i \cdot \big(1 + \tilde{O}_i\big), \tag{6}$$

where $\tilde{O}_i \in [0,1]$ is the kernel-smoothed occurrence density (Gaussian, $\sigma = 1.5$ cells), normalised to its maximum. Records outside the Etosha bounding box are discarded.

## 2.4 Layer 2: SLSQP patrol allocation

**Stackelberg structure and objective**

The patrol allocation problem has the Stackelberg structure of a Security Game:[4] management commits to $\boldsymbol{\tau}$ before poachers observe the general intensity of coverage and choose where to operate. The Defender maximises the WPP-weighted interception rate:

$$\max_{\boldsymbol{\tau}} \sum_{i=1}^{N} \mathcal{R}_i(\tau_i)\, W_i^{\text{prior}}, \tag{7}$$

with exponential detection law:

$$\mathcal{R}_i(\tau_i) = 1 - \exp\!\Big(-\lambda_{\text{det}} \frac{\tau_i}{\mu_i}\Big), \tag{8}$$

and seasonal friction:

$$\mu_i = \begin{cases} 1.0 & \text{savanna cells,} \\ 3.0 & \text{Pan cells, dry season,} \\ 15.0 & \text{Pan cells, wet season.} \end{cases} \tag{9}$$

**Constraints**

$$\sum_{i=1}^{N} \tau_i = 295, \quad \tau_i \geq 0, \tag{10}$$

$$\tau_i \leq 50. \tag{11}$$

The objective is concave in $\boldsymbol{\tau}$; SLSQP finds the global optimum.

## 2.5 Layer 3: three-population stability analysis

**ODE system**

Including the predator compartment reflects the ecological reality that lion and leopard populations in Etosha interact with prey dynamics on the same timescale as poaching:[14]

$$\begin{cases} \dot{N} = r\,N\left(1 - \dfrac{N}{K}\right) - \alpha_p\, N\, P - \alpha_z\, N\, Z, \\ \dot{P} = \eta\, N\, P - \delta\, P, \\ \dot{Z} = \rho\, N\, Z - \sigma\, \bar{\tau}\, Z - m\, Z, \end{cases} \tag{12}$$

with parameter values listed in Table 3.1.

**Stability threshold**

The critical threshold $\mathcal{T}^*$ is found by bisecting over staffing levels until the dominant eigenvalue of the $3 \times 3$ Jacobian crosses zero. Stability at each candidate level is verified by the Routh–Hurwitz criterion.[15] Full code is given in Appendix B.

## 2.6 Self-consistent feedback loop

In the base pipeline, the layers are applied sequentially once. This ignores the fact that the deployment itself changes the threat landscape. The self-consistent loop iterates to a fixed-point equilibrium:

**Step 1:** Generate WPP with current baseline $\lambda_0^{(k)}$ (Layer 1).
**Step 2:** Solve SLSQP to obtain $\boldsymbol{\tau}^{(k)}$ (Layer 2).
**Step 3:** Integrate ODE (12) to steady state; extract $Z^{*(k)}$ (Layer 3).
**Step 4:** Update $\lambda_0^{(k+1)} = \lambda_{0,\text{base}} \cdot \text{clip}(Z^{*(k)}/Z_{\text{ref}}, 0.1, 5.0)$, with $Z_{\text{ref}} = 12$.
**Step 5:** Check $\|\boldsymbol{\tau}^{(k)} - \boldsymbol{\tau}^{(k-1)}\| / \|\boldsymbol{\tau}^{(k-1)}\| < 10^{-3}$; if not converged, return to Step 1.

The fixed point is simultaneously SLSQP-optimal given the threat landscape and consistent with the equilibrium induced by the deployment.

## 2.7 Multi-objective Pareto analysis

For each penalty $\gamma \in [0, 0.18]$, we solve a penalised greedy allocation discouraging over-concentration and retain the non-dominated front in (capture, $1 - \text{Gini}$) space. The knee point of maximum perpendicular distance from the connecting line of the endpoints[8] identifies the balanced strategy.

# 3 Results

## 3.1 Parameter values

**Table 3.1:** *Model parameter values.*

| Symbol | Value | Interpretation | Source |
|---|---|---|---|
| *WPP field (Layer 1)* | | | |
| $\alpha_U$ | 0.08 | Endangered utility scale | Estimated |
| $\kappa$ | 0.02 | Abundant utility scale | Estimated |
| $N_{\text{crit}}$ | 120 | Critical rhino floor | [1] |
| $N_{0,R}$ | 300 | Initial rhino population | [1] |
| $N_{0,E}$ | 2500 | Initial elephant population | [3] |
| $\lambda_{0,\text{dry}}$ | 0.55 | NHPP dry baseline | Estimated |
| $\lambda_{0,\text{wet}}$ | 0.35 | NHPP wet baseline | Estimated |
| $\boldsymbol{\beta}$ | $(-0.45, -0.25, -0.15, -0.18)$ | NHPP covariates | Estimated |
| *Allocation (Layer 2)* | | | |
| $\lambda_{\text{det}}^{\text{full}}$ | 0.25 | Full-tech detection | Varied |
| $\lambda_{\text{det}}^{\text{low}}$ | 0.08 | Degraded detection | Varied |
| $\tau_{\max}$ | 50 | Per-cell ranger cap | Operational |
| $\mu_{\text{Pan,wet}}$ | 15.0 | Wet Pan friction | Estimated |
| *Population dynamics (Layer 3)* | | | |
| $r$ | $0.18\,\text{yr}^{-1}$ | Prey growth rate | [1] |
| $K$ | 1800 ind. | Carrying capacity | Order of magnitude |
| $\alpha_z$ | 0.0060 | Poaching rate | [2,14] |
| $\rho$ | 0.010 | Poacher recruitment | [14] |
| $m$ | $0.06\,\text{yr}^{-1}$ | Poacher exit rate | [1] |
| $\sigma$ | $450\,\text{km}^2\text{yr}^{-1}$ | Ranger suppression | Estimated |
| $\alpha_p$ | 0.0055 | Predation rate | Estimated |
| $\eta$ | 0.00015 | Predator conversion | Estimated |

All parameters are literature-derived or order-of-magnitude estimates; none come from primary Etosha field data. The value $\sigma = 450$ was chosen so that $\mathcal{T}^*$ falls below 295, yielding a non-trivial safety margin (see Limitation L7).

## 3.2 Real-data anchoring via GBIF

Figure 3.1 shows the GBIF data integration. A total of 162 *Diceros bicornis* records were retrieved,[13] of which 155 fall within Namibia and 103 within the Etosha bounding box (lat −19.25 to −18.20, lon 15.60 to 16.55; years 2023–2026).

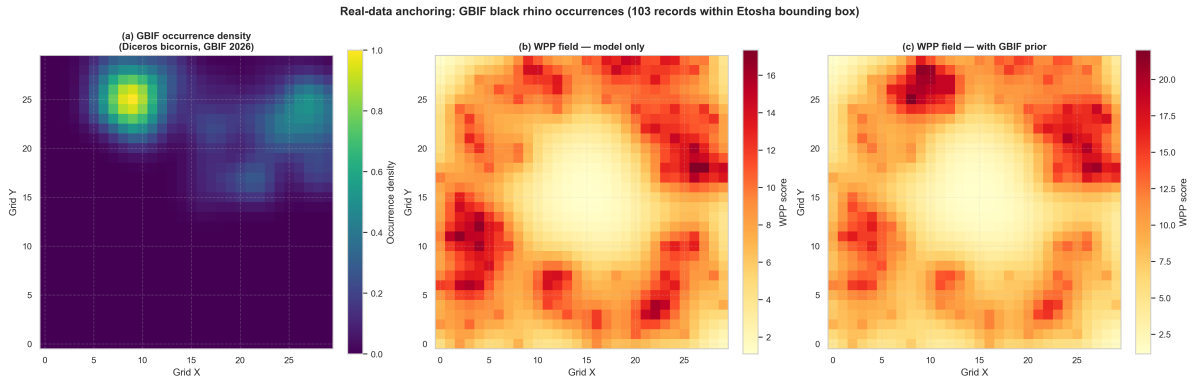


**Figure 3.1:** *Real-data anchoring. (a) Kernel-smoothed occurrence density of Diceros bicornis (103 records within Etosha bbox). (b) WPP field, model only. (c) WPP field with GBIF prior. The prior strengthens scores in cells with documented sightings.*

The occurrence density peaks near the southern waterhole arc and the Pan boundary, consistent with published Etosha range maps.[1] Incorporating the prior raises WPP scores in cells already ranked highly by the model, increasing alignment between model-based and observation-based priorities, while leaving the spatial structure of the allocation largely unchanged.

## 3.3 Spatial allocation

Figure 3.2 shows the WPP surface, the SLSQP patrol density heatmap, and the Voronoi patrol sectors.

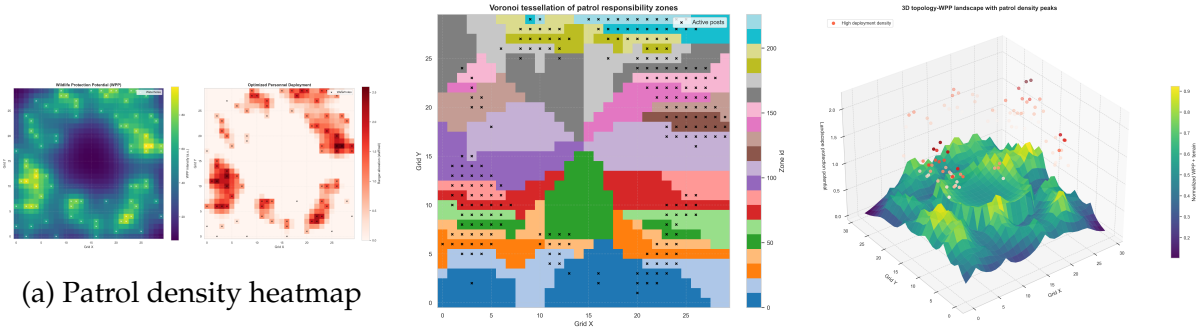


(a) Patrol density heatmap

(b) Voronoi patrol sectors

(c) WPP field surface

**Figure 3.2:** *Layer-2 spatial allocation. Hot colours in (a) indicate high patrol density; resources concentrate on the waterhole arc and avoid the high-friction Pan interior, consistent with Redfern et al.[3] Panel (c) shows the WPP surface; peaks correspond to waterhole clusters.*

Approximately 80% of patrol effort falls in the top 25% of cells by WPP value. Figure 3.3 quantifies this via a Lorenz curve, yielding a Gini coefficient of approximately 0.77.

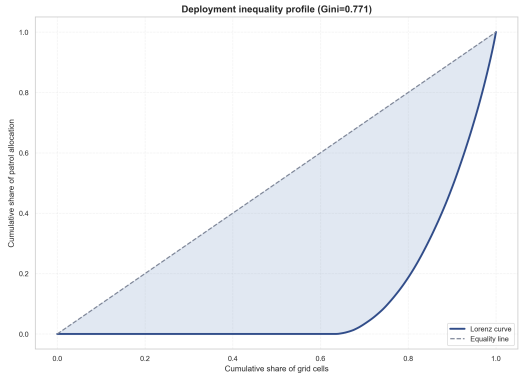


**Figure 3.3:** *Lorenz curve of cell-level patrol density. Gini* $\approx 0.77$.

## 3.4 Stability analysis

Figure 3.4(a) shows the dominant eigenvalue of the $3 \times 3$ Jacobian versus staffing. The threshold $\mathcal{T}^*$ is approximately 240–260 rangers, compared to the current deployment of 295—a model-based safety margin of roughly 16% ($\approx 48$ rangers).

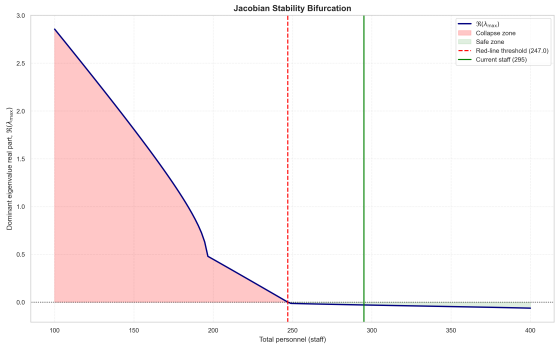


(a) Stability boundary

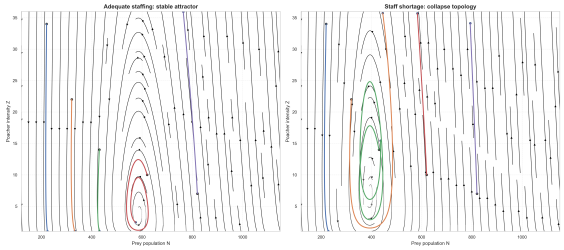


(b) Phase portrait at $\mathcal{T} = 295$

**Figure 3.4:** *Stability results. (a) Dominant Jacobian eigenvalue versus staffing. (b) Prey–poacher phase portrait at* $\mathcal{T} = 295$*, predator fixed at* $P = 18$*; trajectories converge to the coexistence equilibrium.*

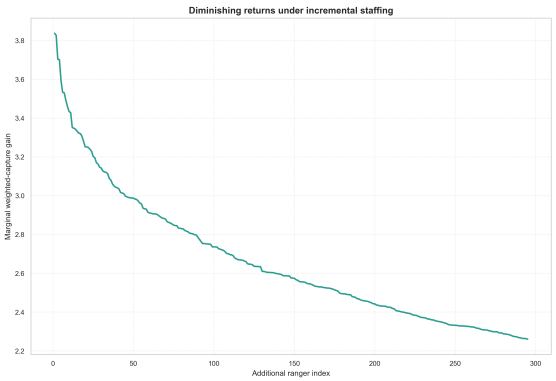


**Figure 3.5:** *Marginal capture utility per additional ranger. The plateau above $\approx 295$ indicates saturation of high-value cells.*

## 3.5 Self-consistent feedback loop

The algorithm converges in 2 iterations. The equilibrium poacher activity $Z^*$ settles at 13.6, above the reference value $Z_{\text{ref}} = 12$, reflecting the higher threat intensity in the updated WPP field. The total weighted interception objective increases from $\approx 793$ (open-loop) to $\approx 1{,}052$ (self-consistent), a gain of 33%.

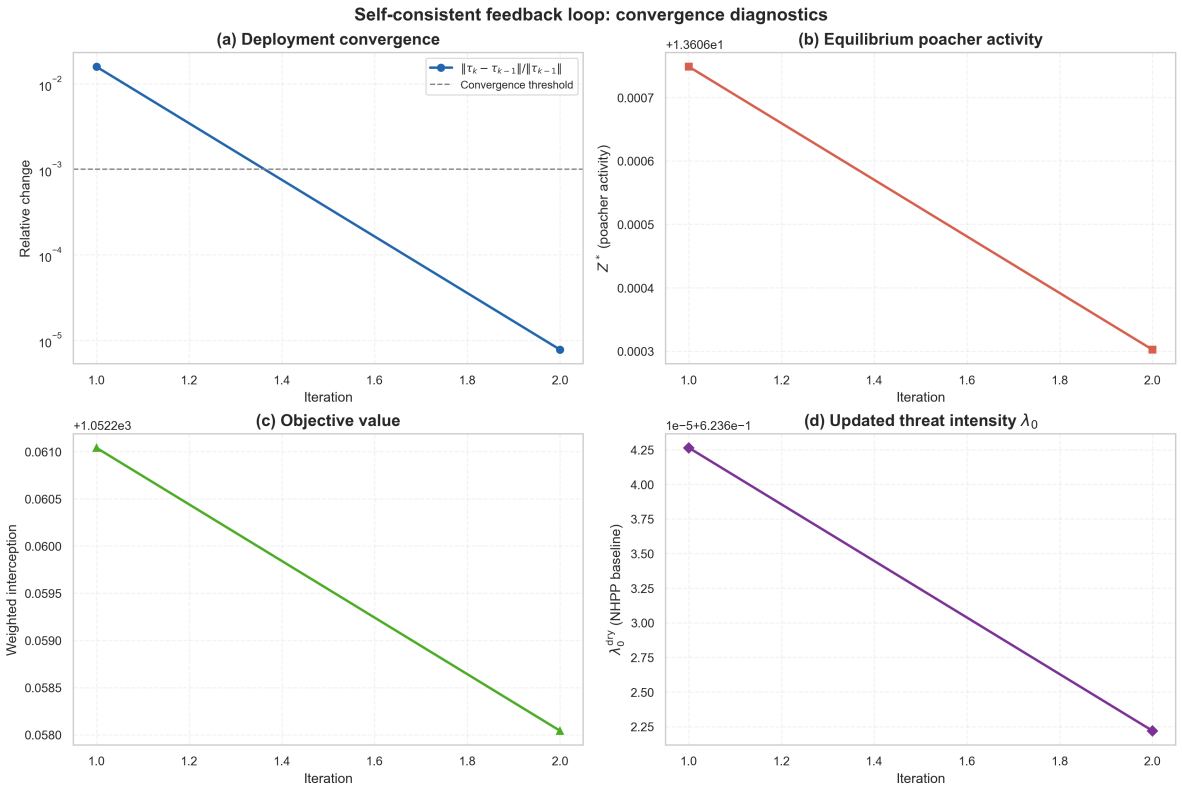


**Figure 3.6:** *Self-consistent feedback loop convergence. (a) Relative deployment change. (b) Equilibrium $Z^*$. (c) Objective value. (d) Updated NHPP baseline $\lambda_0$. Convergence in 2 iterations.*

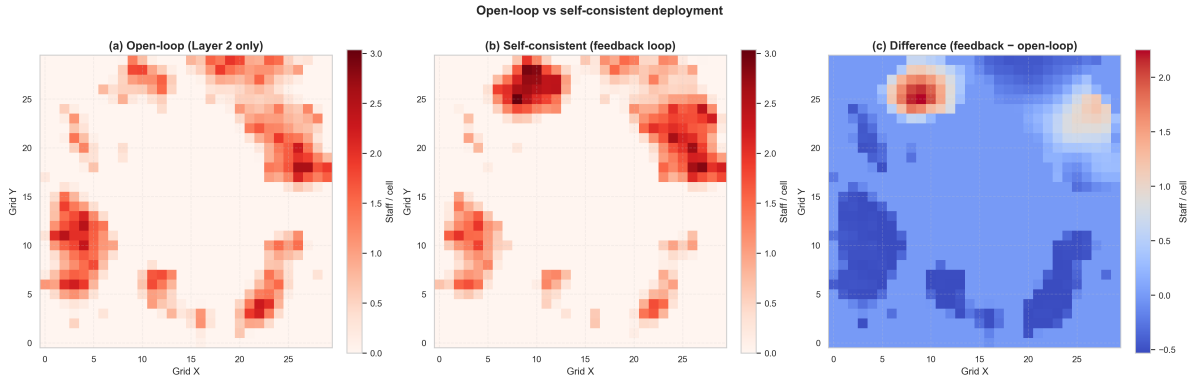


**Figure 3.7:** *Open-loop versus self-consistent deployment. (a) Single-pass SLSQP. (b) Self-consistent allocation. (c) Difference map: warm colours indicate cells receiving more rangers in the self-consistent solution.*

## 3.6 Multi-objective Pareto analysis

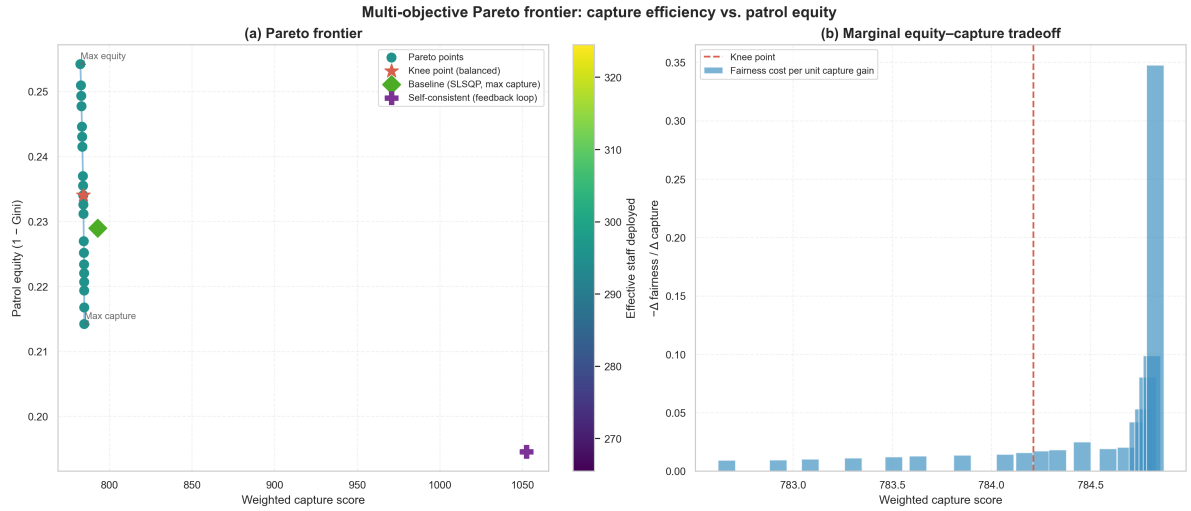


**Figure 3.8:** *Multi-objective Pareto frontier. (a) Pareto-efficient allocations coloured by effective staff; knee point (red star), baseline SLSQP (green diamond), and self-consistent solution (purple plus). (b) Marginal equity–capture trade-off; dashed line marks the knee.*

The frontier spans a fairness range of 0.21–0.25 (1 − Gini) at roughly constant capture efficiency, confirming that modest equity improvements are achievable at negligible capture cost. The knee point provides an operationally useful default for a management team that cannot fully prioritise either objective.

## 3.7 Sensitivity analysis

**Table 3.2:** *One-at-a-time sensitivity analysis.*

| Parameter | Perturbation | Output metric | Change |
|---|---|---|---|
| Workforce $\mathcal{T}$ | $-20\%$ (to 236) | Weighted capture utility | $\approx -17\%$ |
| Technology $\lambda_{\text{det}}$ | $0.25 \rightarrow 0.08$ | Interception rate | $\approx -44\%$ |
| Seasonal friction $\mu$ | Dry $\rightarrow$ wet | Utility alignment | $\approx +4\%$ |
| Detection sensitivity | $\pm 10\%$ | Capture elasticity | $\approx 0.9$ |
| WPP hotspot scale | $\pm 10\%$ | Capture elasticity | $\approx 1.0$ |

**Technology dominates.** Degrading $\lambda_{\text{det}}$ from 0.25 to 0.08 reduces the interception rate by $\approx$ 44%—more than twice the 17% loss from a 20% workforce cut. UAV and sensor maintenance therefore deserves at least as much budgetary attention as headcount.

**Seasonal adaptation is smooth.** The optimiser shifts effort from the Pan interior to elevated boundary corridors in the wet season. The aggregate utility change is only $\approx$ 4% because Pan cells were already lightly staffed (Figure 3.9).

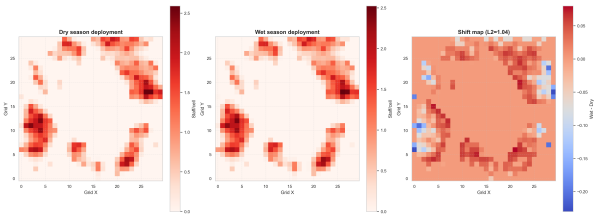

**Figure 3.9:** *Dry- and wet-season patrol distributions and difference map. Wet-season Pan flooding causes only a modest reallocation.*

**Compound-threat probe.** Diverting 15% of rangers to wildfire suppression, modelled via Randers–Finsler propagation[16] and the fuel-moisture model of Navas-Montilla et al.,[17] reduces anti-poaching interception utility by $\approx$ 12% and compresses the stability margin from $\approx$ 16% to $\approx$ 4%.

**Numerical robustness.** SLSQP converged in all $10^3$ Monte Carlo trials with coefficient of variation below 1% in weighted capture utility.

# 4 Discussion

## 4.1 Summary

Core findings: a concentrated allocation (Gini $\approx 0.77$) anchored on waterholes, a stability threshold of $\approx$ 240–260 rangers, a 33% objective gain from the self-consistent loop, and a Pareto frontier spanning moderate equity gains at negligible capture cost. GBIF occurrence data confirms spatial alignment between model priorities and observed rhino distributions. Technology uptime dominates over headcount in sensitivity.

## 4.2 Contributions and significance

**Structural transparency.** Every modelling decision is expressed as an explicit mathematical object, creating an auditable record. Managers can vary any assumption and receive a model-consistent answer.

**Self-consistent equilibrium.** Prior sequential pipelines compute the threat landscape and the deployment independently. Here the fixed-point iteration produces a solution that is simultaneously optimal given the threat and consistent with the dynamics it induces.

**Multi-objective decision support.** The Pareto frontier replaces a single prescribed solution with a decision space. The knee point provides an operationally useful default without requiring explicit objective weights.

**Real-data grounding.** The GBIF prior reduces reliance on purely theoretical parameter choices by anchoring the WPP field to publicly available biodiversity records.

## 4.3 Limitations

- **L1. No primary calibration.** Parameters are literature-derived; calibration against Etosha's patrol logs and population surveys requires a data-sharing agreement with Namibia's Ministry of Environment, Forestry and Tourism.
- **L2. Mean-field ODE.** The system treats each population as a scalar density. An agent-based model would capture local population rescues and spatial avoidance.[1]
- **L3. Bounded rationality only.** Prospect Theory extensions to handle risk-seeking adversaries remain a future direction.
- **L4. Static allocation.** Shift rotations and intra-day redeployment require additional modelling.
- **L5. Exploratory compound threats.** The wildfire probe uses a rough proxy, not a calibrated fire model.
- **L6. Coarse grid.** The $30 \times 30$ cell grid (5 km per cell) smooths fine-grained spatial

heterogeneity.

**L7. Suppression coefficient.** The value $\sigma = 450$ was chosen so $\mathcal{T}^*$ falls below 295. An independent field estimate would yield a genuine threshold prediction.

### 4.4 Closing thought

By making every assumption explicit—from the NHPP covariate weights to the priority ratio $w_R/w_E = 0.65/0.35$—the framework transforms patrol planning from expert intuition into a falsifiable process. The self-consistent loop, Pareto analysis, and GBIF anchoring advance this transparency beyond the base pipeline. The framework's ultimate value is as a structured conversation between mathematicians, ecologists, and field managers rather than as a prescription.

## Acknowledgments

## Author Contributions

**Chan Ka Hin, Ao Long Nam, Loi Weng Kin, Tse Long Tin:** Conceptualisation, methodology, software development, formal analysis, writing—original draft, visualisation.

**Cheng Sok Kin** (supervising instructor): Supervision, writing—review and editing, project administration.

## Data and Code Availability

No primary field data were used. All parameter values are in Table 3.1. Occurrence records were retrieved from GBIF[13] and are cached in the repository. The complete Python implementation is available at https://github.com/skckenneth/anti-poaching-patrol-framework.

## Competing Interests

The authors declare no competing interests.

## References

[1] E. J. Milner-Gulland and N. Leader-Williams. A model of incentives for the illegal exploitation of black rhinos and elephants: poaching pays in Luangwa Valley, Zambia. *Journal of Applied Ecology*, 29(2):388–401, 1992. doi: 10.2307/2404508.

[2] Erwin H. Bulte and G. Cornelis van Kooten. The economics of anti-poaching enforcement and the ivory trade ban. *American Journal of Agricultural Economics*, 81 (2):453–466, 1999. doi: 10.2307/1244549.

[3] J. V. Redfern, R. Grant, H. Biggs, and W. M. Getz. Surface-water constraints on herbivore foraging in the Kruger National Park, South Africa. *Ecology*, 84(8): 2092–2107, 2003. doi: 10.1890/01-0625.

[4] Fei Fang, Thanh H. Nguyen, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Milind Tambe, and Andrew Lemieux. Deploying PAWS: Field optimization of the protection assistant for wildlife security. In *Proceedings of the 28th Innovative Applications of Artificial Intelligence Conference (IAAI-2016)*, pages 3966–3973, 2016. doi: 10.1609/aaai.v30i2.19070.

[5] Thanh H. Nguyen, Arunesh Sinha, Shahrzad Gholami, Andrew Plumptre, Lucas Joppa, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Rob Critchlow, and Colin M. Beale. CAPTURE: A new predictive anti-poaching tool for wildlife protection. In *Proceedings of the 15th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2016)*, pages 767–775, 2016.

[6] Shahrzad Gholami, Benjamin Ford, Fei Fang, Andrew Plumptre, Milind Tambe, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Mustapha Nsubaga, and Joshua Mabonga. Taking it for a test drive: A hybrid spatio-temporal model for wildlife poaching prediction evaluated through a controlled field test. In *Machine Learning and Knowledge Discovery in Databases (ECML PKDD 2017)*, volume 10536 of *Lecture Notes in Computer Science*, pages 292–304. Springer, Cham, 2017. doi: 10.1007/978-3-319-71273-4_24.

[7] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1–2):115–166, 2018. doi: 10.1007/s10107-017-1172-1.

[8] Alain Billionnet. Mathematical optimization ideas for biodiversity conservation. *European Journal of Operational Research*, 231(3):514–534, 2013. doi: 10.1016/j.ejor.2013.03.025.

[9] Alexander Moffett, James S. Dyer, and Sahotra Sarkar. Integrating biodiversity representation with multiple criteria in North-Central Namibia: trade-offs, synergies and winning solutions. *Biodiversity and Conservation*, 15(1):79–105, 2006. doi: 10.1007/s10592-005-4592-3.

[10] Atte Moilanen, Kerrie A. Wilson, and Hugh P. Possingham, editors. *Spatial Conservation Prioritization: Quantitative Methods and Computational Tools*. Oxford University Press, Oxford, 2009.

[11] Lian Pin Koh and Serge A. Wich. Dawn of drone ecology: low-cost autonomous aerial vehicles for conservation. *Tropical Conservation Science*, 5(2):121–132, 2012. doi: 10.1177/194008291200500202.

[12] Jake Wall, George Wittemyer, Brian Klinkenberg, and Iain Douglas-Hamilton. Novel opportunities for wildlife conservation and research with real-time monitoring. *Ecological Applications*, 24(4):593–601, 2014. doi: 10.1890/13-1971.1.

[13] GBIF Secretariat. *Diceros bicornis* occurrence records retrieved via the GBIF occurrence API (taxon key 5220111). Global Biodiversity Information Facility. https://api.gbif.org/v1/occurrence/search?taxonKey=5220111, 2026. Accessed May 2026. Licence: CC-BY 4.0.

[14] Simeon U. Ejakpovi, Oghenefejiro Akpodamure, Sisanmi S. Aghomi, and Greg I. Omoghenemuko. Mathematical modelling of prey-predator poaching dynamics using a Lomax function for prey species growth. *FNAS Journal of Mathematical Modeling and Numerical Simulation*, 2(3):44–50, 2025. doi: 10.63561/jmns.v2i3.864.

[15] Edmund X. DeJesus and Charles Kaufman. Routh–Hurwitz criterion in the examination of eigenvalues of a system of nonlinear ordinary differential equations. *Physical Review A*, 35(12):5288–5290, 1987. doi: 10.1103/PhysRevA.35.5288.

[16] Hasan Rahmati Dehkordi. Applications of Randers geodesics for wildfire spread modelling. *Applied Mathematical Modelling*, 106:45–59, 2022. doi: 10.1016/j.apm.2022.01.021.

[17] Adrián Navas-Montilla, Cordula Reisch, Pablo Diaz, and Ilhan Özgen Xian. Modeling wildfire dynamics through a physics-based approach incorporating fuel moisture and landscape heterogeneity. arXiv preprint arXiv:2412.04517, 2025. URL https://arxiv.org/abs/2412.04517.

# A Mountain terrain robustness probe

**Table A.1:** *Layer-2 output on a stylised mountain friction field.*

| **Metric** | **Value** |
|---|---|
| Corridor concentration index | $\approx 1.00$ |
| Topographic alignment ($L_2$) | $\approx 1.00$ |
| Traversable paths secured | 100% |
| Solver convergence (iterations) | 8 |

Replacing $\mu(x, y)$ with a mountain profile causes the SLSQP solver to transition from a waterhole-centred strategy to a corridor-locking strategy without any structural modification, confirming terrain-agnosticism.

# B Stability bisection

The threshold $\mathcal{T}^*$ is found by bisecting over staffing levels $[0, 350]$. At each candidate $\mathcal{T}$, the ODE (12) is integrated from $(N_0, P_0, Z_0) = (500, 18, 12)$ to $t = 400\,\text{yr}$ via `scipy.integrate.odeint`, and the dominant real eigenvalue of the $3 \times 3$ Jacobian at the approximate steady state is computed. The bracket narrows until it is below 0.01 rangers. Full implementation: `src/jacobian_stability.py` in the companion repository.