Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A Spatial Model of Chronic Wasting Disease and the Optimal Placement of a Sight-Based Cull

Gabe Flanagin

gmflanagin@gmail.com

Western Washington University https://orcid.org/0009-0004-9113-4369

Alex Escoto

Western Washington University

Zhen Chao

Western Washington University

Research Article

Keywords: Chronic wasting disease, reaction-diffusion, PDE-constrained optimal control, discrete adjoint, surface finite elements, landscape heterogeneity,

Posted Date: November 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11598596/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

A Spatial Model of Chronic Wasting Disease and the Optimal Placement of a Sight-Based Cull

Gabe Flanagin* Alex Escoto† Zhen Chao‡

Abstract

Chronic wasting disease (CWD) is a fatal prion disease of cervids that spreads by direct contact and through a persistent environmental reservoir. We ask where and when a cull restricted to the clinically affected animals that field crews can identify should be placed on a real landscape, and what it can achieve. We formulate a five-compartment (susceptible, exposed, pre-clinical, clinical, environmental) reaction–diffusion model on a surface built from digital elevation data, with land cover setting the carrying capacity and slope and land cover setting an anisotropic diffusion tensor. We optimize a space- and time-dependent removal rate on the clinical class, computing gradients with the exact discrete adjoint of the finite element scheme. A closed-form bound gives the lowest control reproduction number such a cull can reach. With our parameters, pre-clinical animals generate 39.5% of local transmission, and maximal effort leaves it at or above 1.36 across our sweep. Simulations across reservoir persistence and transmissibility confirm that the optimal cull slows but never halts the invasion. The front retains 56–94% of its uncontrolled speed, and the control recovers 7–53% of the accumulated loss of susceptible deer. Where the herd is depleted, optimal effort forms a ring at the leading edge, 12–19 km ahead of the clinical peak. Cumulative incidence can rise under control even as per-capita risk falls, because the cull keeps more susceptible deer alive. On three Colorado landscapes, terrain and land cover barely change the benefit, and the cull needs 26–35% less effort, since less ground carries the invasion.

Keywords: chronic wasting disease; reaction–diffusion; PDE-constrained optimal control; discrete adjoint; surface finite elements; landscape heterogeneity

Mathematics Subject Classification: 92D30; 49K20; 35K57; 65M60

1 Introduction

Chronic wasting disease (CWD) is a fatal, transmissible prion disease of deer, elk, moose and other cervids, and it has now been detected across much of North America as well as in Scandinavia and South Korea (Escobar et al., 2020). Unlike most wildlife epizootics, it spreads by two routes. Infected animals shed prions in saliva, urine and feces long before they show clinical signs (Mathiason et al., 2006), so transmission occurs both through direct contact and through a contaminated environment. Prions bind to soil minerals and remain infectious (Johnson et al., 2006), and paddocks that once held infected deer transmitted the disease to naive animals two years after the infected animals were removed (Miller et al., 2004). Since the reservoir persists independently of

*Corresponding author. Department of Mathematics, Western Washington University, Bellingham, WA, USA. flanagg@wwu.edu

†Department of Mathematics, Western Washington University, Bellingham, WA, USA. escota@wwu.edu

‡Department of Mathematics, Western Washington University, Bellingham, WA, USA. chaoz2@wwu.edu

its hosts, CWD can persist at host densities too low to sustain direct transmission alone (Almberg et al., 2011). Field studies now link endemic CWD to population declines in both white-tailed deer (Edmunds et al., 2016) and mule deer (DeVivo et al., 2017). There is no vaccine or treatment, and management consists largely of removing deer, through harvest or targeted culling, with little empirical evidence on how that effort should be placed (Uehlinger et al., 2016).

Mathematical models have shaped much of what is known about CWD dynamics. Well-mixed compartmental models established the roles of direct and environmental transmission (Miller et al., 2006; Almberg et al., 2011), showed that environmental persistence can drive population decline and extinction (Almberg et al., 2011) and recurrent epizootics (Sharp and Pastor, 2011), and examined how harvest and predation interact with transmission (Potapov et al., 2016; FitzGerald and Keener, 2025). Management-oriented models have compared culling strategies (Wasserberg et al., 2009) and contrasted selective removal of infected animals, as by predators, with non-selective removal (Wild et al., 2011; Miller et al., 2008). Spatial models have followed more slowly. Ecological diffusion models fitted to telemetry and surveillance data reproduce observed spread in mule deer in Utah (Garlick et al., 2014) and forecast spread in white-tailed deer in Wisconsin (Hefley et al., 2017), and a recent causal analysis of targeted culling in Wisconsin found that it reduced the spatial diffusion of the disease rather than its local growth rate (Mandujano Reyes et al., 2025).

Two recent studies show where the gap lies. Suresh et al. (2026) fit a susceptible–infected–environmental system on a 20 km grid over Kansas, coupled between cells by a mixing matrix built from the order of first detection, and recover the direction and rate of spread from nineteen years of surveillance. Their model has no management lever, and the authors identify the absence of land cover, soil and elevation inputs as its main limitation. Yin et al. (2026) pose a stochastic, well-mixed susceptible–infected–environmental system and use deep reinforcement learning to find a control policy over hunting and environmental decontamination, and name spatial structure as a natural extension. Neither model can ask where on a real landscape culling effort should go.

Optimal control theory offers a principled way to ask that question. It has a long history in the management of wildlife disease (Lenhart and Workman, 2007; Clayton et al., 2010), and Miller Neilan and Lenhart (2011) extended it to a parabolic epidemic system to optimize the spatial distribution of rabies vaccine baits. Modern PDE-constrained optimization provides the adjoint machinery that makes such problems tractable on realistic meshes (Hinze et al., 2009). To our knowledge, these tools have not been applied to CWD, whose environmental reservoir, long pre-clinical stage and strong dependence on habitat make spatial placement both difficult and important.

In this paper we make the CWD model spatially explicit and continuous, with terrain and land cover entering both the carrying capacity and the diffusion tensor, and optimize a culling effort that varies over space and time. Throughout, the cull is restricted to the one class a field crew can identify without a diagnostic test, the clinically affected animals. Our contributions are as follows.

1. We formulate a susceptible–exposed–pre-clinical–clinical–environmental (S-E-I-R-W) reaction–diffusion model on a surface $\Gamma \subset \mathbb{R}^3$ built from digital elevation data, with an anisotropic diffusion tensor that redirects movement along terrain contours and with every parameter traced to its source (Section 2).

2. We derive a closed-form, model-independent bound on what a sight-based cull can achieve, expressed through R_0 and three ratios that can in principle be measured in the field (Section 3.2).

3. We state the optimal control problem with a sparsity-promoting cost, derive its adjoint system and optimality condition, and implement the exact discrete adjoint of a surface finite element

scheme, verified by Taylor tests to second order (Sections 3–6).

4. Through a factorial sweep on a flat domain and an ablation on three Colorado landscapes, we characterize where optimal effort falls relative to the invasion front, what it achieves, and what terrain and land cover change (Section 7).

Section 8 discusses the management implications and limitations, and Section 9 concludes. Each parameter value is given where the parameter first appears, and the values are collected in Table 1 for the model and in Section 3.3 for the control problem.

2 The Reaction–Diffusion Model on a Terrain Surface

2.1 The domain: an embedded surface Γ

The domain $\Gamma \subset \mathbb{R}^3$ is a two-dimensional surface interpolated from the digital elevation model (DEM) of the U.S. Geological Survey (USGS) 3D Elevation Program (U.S. Geological Survey, 2024a), resampled to a 30 m grid, and triangulated into a mesh Γ_h (Section 6) whose vertices carry a position $\mathbf{x} \in \mathbb{R}^3$ and whose triangles carry a unit surface normal $\mathbf{n}(\mathbf{x})$. Land cover from the Annual National Land Cover Database (NLCD; U.S. Geological Survey, 2024b) is aligned to the surface by longitude and latitude, so each point of Γ is tagged with a land-cover class that sets both K and $\mathbf{D}$. All spatial operators are surface operators: $\nabla_\Gamma u$ is the tangential gradient of a scalar u and $\nabla_\Gamma \cdot \mathbf{F}$ the surface divergence of a tangential field $\mathbf{F}$.

2.2 State variables and governing equations

We split the deer population into four host densities (deer/km^2): susceptible S; exposed E, which are infected but latent and neither transmit nor shed; pre-clinical I, which are infectious and shedding but look healthy in the field; and clinical R, which are visibly affected, shed faster and move much less. We write $N = S + E + I + R$ for the total density. A fifth field W is the environmental prion load deposited through saliva, urine, feces and carcasses. These are the stages of Almberg et al. (2011), resolved one for one, with removal at rate α representing death. Separating I from R lets us represent a *sight-based* cull (Section 3), since R is the set of animals a field crew can identify without a diagnostic test.

The governing system on Γ, with a dot denoting ∂_t, is

$$\dot{S} = \nabla_\Gamma \cdot \big(\mathbf{D}_S \nabla_\Gamma S\big) + rS\Big(1 - \frac{N}{K}\Big) - S\,\Lambda, \tag{1}$$

$$\dot{E} = \nabla_\Gamma \cdot \big(\mathbf{D}_E \nabla_\Gamma E\big) + S\,\Lambda - \sigma_1 E, \tag{2}$$

$$\dot{I} = \nabla_\Gamma \cdot \big(\mathbf{D}_I \nabla_\Gamma I\big) + \sigma_1 E - \sigma_2 I, \tag{3}$$

$$\dot{R} = \nabla_\Gamma \cdot \big(\mathbf{D}_R \nabla_\Gamma R\big) + \sigma_2 I - \alpha R, \tag{4}$$

$$\dot{W} = \rho_I I + \rho_R R - \delta W, \tag{5}$$

where the force of infection per susceptible combines the two transmission routes,

$$\Lambda \;:=\; \beta_1\big(I + R\big) \;+\; \beta_2\, W. \tag{6}$$

Both shedding stages transmit directly at the same rate β_1 (Section 2.7). We specify the tensors $\mathbf{D}_X$ in Section 2.5 and the carrying capacity $K(\mathbf{x})$ in Section 2.4. The system is closed with no-flux conditions

$$\big(\mathbf{D}_X \nabla_\Gamma X\big) \cdot \boldsymbol{\nu} = 0 \qquad \text{on } \partial\Gamma, \qquad X \in \{S, E, I, R\}, \tag{7}$$

where $\boldsymbol{\nu}$ is the outward conormal, so deer neither enter nor leave the domain.

2.3 Reaction kinetics

Recruitment $rS(1-N/K)$ is logistic. Density dependence acts through the total density N, and all newborns enter S. Direct transmission $\beta_1 S(I+R)$ comes from contact with shedding deer of either stage, and environmental transmission $\beta_2 SW$ from contact with contaminated ground. Infected animals leave the latent and pre-clinical stages at rates σ_1 and σ_2, and αR is the removal of clinical deer by disease-induced mortality, together with any predation or harvest that selectively takes symptomatic animals. There is no recovery, so every exposed animal eventually becomes clinical. Deposition $\rho_I I + \rho_R R$ combines shedding by living animals with the release at death and during carcass decomposition, and δW is the loss of infectivity through degradation, burial or transport, with $1/\delta$ the mean persistence time.

Splitting the shedding period into two stages has two effects. First, the time spent shedding is Erlang rather than exponentially distributed, which changes the generation-time distribution and with it the growth rate at fixed R_0. Second, it lets the stages differ in mobility, since deer slow down once clinical signs appear. The two transmission routes also act on different timescales. Direct transmission follows the current shedding population, whereas W integrates past shedding over a memory of $1/\delta$, so present risk is decoupled from present prevalence. We treat prions as immobile, so (5) is a local ordinary differential equation (ODE) at each point of Γ.

2.4 Carrying capacity K and demography

We draw parameter values from the CWD literature, using mule deer data wherever possible, and label each value by its provenance: [LIT], derived from published CWD data; [DEF], a unit-defining normalization; [CAL], back-solved from a stated calibration target; [ASM], assumed, that is, expert judgment without a supporting citation for this region; and [NUM], a conditioning or discretization parameter that represents nothing about deer or prions. Table 1 collects them, and Section S1 of Online Resource 1 records how each value was reached.

The carrying capacity is a base capacity K_0 (deer/km^2) scaled by a land-cover coefficient,

$$K(\mathbf{x}) = K_0\, c(\mathbf{x}), \tag{8}$$

where $c(\mathbf{x}) \in [0,1]$ is the habitat quality of the local land-cover class. We take c from the mule deer motilities of Garlick et al. (2014) for the five classes they resolve and set the rest by expert judgment (Table 2 and Section 2.5). This makes K discontinuous across class boundaries, and the discretization smooths each boundary over a single element (Section 6). We take $K_0 = 3$ deer/km^2 in the best habitat class, [ASM]. Three independent readings of deer density (Geremia et al., 2015; Miller et al., 2006; Suresh et al., 2026) bracket it to 2–3, and Section S1 of Online Resource 1 reconciles them.

The intrinsic growth rate is $r = \ln 1.30 = 0.26$ yr^{-1}, [LIT], the maximum annual finite growth rate that Almberg et al. (2011) report for mule deer. This is already the density-free rate the logistic term requires. The life-table eigenvalues of Tanner (1975) give at most 0.24. Since a larger r speeds recovery, our value errs in the control's favor (Section S1 of Online Resource 1).

2.5 Terrain- and land-cover-aware diffusion tensor D

Climbing is energetically costly for mule deer (Parker et al., 1984), and mountain ungulates cross slopes obliquely, so the grade they climb is shallower than the terrain (Redcliffe et al., 2025). We therefore let slope redirect movement onto the contour rather than remove it, and scale both

directions by the permeability of the local land cover. With $\mathbf{e}_\mathrm{s}$ and $\mathbf{e}_\mathrm{c}$ the unit tangent vectors of Γ along the slope and along the contour, the diffusion tensor is

$$\mathbf{D}(\mathbf{x}) = \kappa\, p(\mathbf{x}) \Big[\eta(\theta)\, \mathbf{e}_\mathrm{s} \otimes \mathbf{e}_\mathrm{s} + \big(2 - \eta(\theta)\big)\, \mathbf{e}_\mathrm{c} \otimes \mathbf{e}_\mathrm{c} \Big], \tag{9}$$

where $p(\mathbf{x}) \in (0, 1]$ is the permeability of the local land-cover class, set below. The scale κ carries the physical magnitude of the diffusivity, and every other factor in (9) is dimensionless. Since $\mathbf{D}$ is symmetric, the surface diffusion operator is self-adjoint. The trace of $\mathbf{D}$ is $2\kappa p$ at every slope, as for a walker with fixed speed u and persistence time t_p, whose $\mathbf{D} = u^2 t_\mathrm{p} \langle \mathbf{e} \otimes \mathbf{e} \rangle$, averaged over its headings $\mathbf{e}$, has a trace that does not depend on how those headings are distributed. Slope thus changes the direction in which deer travel but not how far. The factor η sets how strongly movement along the fall line is suppressed, and the fixed trace sends that movement onto the contour. Deer on steep ground likely also move less overall, but we do not include this, since it would add a parameter with no data to set it.

Along the slope, diffusion is suppressed by the factor

$$\eta(\theta) = \iota + (1 - \iota)\, \mathrm{e}^{-\gamma \tan\theta}, \tag{10}$$

where θ is the angle between $\mathbf{n}$ and the vertical, so $\tan\theta$ is the grade, which is finite everywhere since Γ is the graph of an elevation function. The exponential form follows Tobler's hiking function, in which walking speed falls exponentially with grade (Tobler, 1993), so the penalty begins at the first degree of slope and is set by a single rate γ. On flat ground $\eta = 1$ and the tensor is isotropic, so it does not depend on $\mathbf{e}_\mathrm{s}$ where the slope direction is undefined, and the floor $\iota > 0$ keeps it nonsingular along $\mathbf{e}_\mathrm{s}$. Each host compartment also carries a mobility multiplier, $\mathbf{D}_X = D_X^\mathrm{scale} \mathbf{D}(\mathbf{x})$. Susceptible, exposed and pre-clinical animals move alike, since neither latency nor pre-clinical shedding is visible in the field, while clinical animals move less (Section 2.6).

Since no diffusion coefficient has been fitted for mule deer, we fix κ by matching an observed spread rate. Every host diffusivity is proportional to κ, so front speeds scale as $\sqrt{\kappa}$, and one model speed converts any observed speed into a required κ. We adopt $\kappa = 4 \times 10^6$ m^2/yr $= 4$ km^2/yr on flat, fully permeable cover, [ASM] in value and [LIT] in the rate it matches. With this value the uncontrolled leading edge (Section 7.2) advances at 0.6–2.9 km/yr across the sweep, and at 1.6 km/yr at ($\delta = 0.25$, $R_0 = 3.5$), the cell shared with the landscape runs. The single-species Fisher–KPP speed $2\sqrt{\kappa r_\mathrm{inv}}$, with r_inv the invasion growth rate of the linearized infection, overstates this, since W does not move and clinical animals move at 0.7 of the base rate. Reported CWD spread rates range from 1.6 to 5.5 km/yr. At the shared cell, $\kappa = 4$ matches the northwestern Kansas zone of Suresh et al. (2026), while the 3.7–5.5 km/yr of Garlick et al. (2014) would need $\kappa \approx 21$–47. We choose the slow end because the domain does not scale with κ. A diffusivity an order of magnitude larger carries the invasion across any of our landscapes within a few years at the higher reproductive numbers, leaving boundary effects rather than a front for the control to act on. Even at $\kappa = 4$, the uncontrolled invasion reaches the rim of the 40 km sweep disk within the control horizon at eight of twelve cells, so the diffusivity and the domain size are one design choice. On the flat disk of Section 7.1 this choice costs nothing, since there κ only sets the length scale. Every time and percentage holds at any κ', while distances and speeds scale by $\sqrt{\kappa'/\kappa}$. Our 40 km disk at $\kappa = 4$ is thus also a 137 km disk at $\kappa = 47$. The landscapes do not rescale this way, since terrain and land cover do not stretch with κ (Section S2 of Online Resource 1).

We take the rate $\gamma = 4.7$ from mule deer energetics, [LIT]. Parker et al. (1984) measured the net cost of travel on a 14.3° incline, and for a 66.5 kg adult the cost averaged over ascent and descent is 1.83 times the level cost. At fixed power, speed is inversely proportional to cost, so with $\mathbf{D} \propto u^2$ we have $\eta(14.3°) = 1.83^{-2} = 0.30$, which gives $\gamma = 2 \ln 1.83 / \tan 14.3° = 4.7$. The factor falls to

0.07 at 30° and 0.01 at 45°. If diffusivity is instead taken proportional to speed, as in a resistance surface, γ halves to 2.4. The floor $\iota = 10^{-3}$ is a conditioning parameter rather than a movement parameter, [NUM]. The exponential reaches it only above 56°, so it binds on at most 0.02% of the elements of any mesh, and it sets the finest edge of the mesh (Section 6). Mesh coordinates are in meters and densities per km^2. Since mesh length enters only through diffusion, κ is the one parameter given in m^2/yr.

The land-cover coefficients rest on the form of the diffusion term. We write it in divergence form, $\nabla_\Gamma \cdot (\mathbf{D}\nabla_\Gamma X)$, under which p sets how readily deer cross a class and, without reaction, a uniform density stays at rest. Under the ecological diffusion $\Delta(\mu X)$ of Garlick et al. (2014), deer instead accumulate where their motility μ is low, and the steady state $X \propto 1/\mu$ measures habitat use rather than movement. A published coefficient can therefore supply p only if it was fitted in divergence form, while an ecological motility measures what c does. We take p from the divergence-form Fisher–KPP fit of Mandujano Reyes et al. (2025) to CWD surveillance of white-tailed deer in Wisconsin, and c from the inverse winter motilities that Garlick et al. fit to mule deer GPS collars. The two need not agree, since barren land is easy to cross but supports few deer and pasture is the reverse. Classes that neither source resolves are expert judgment, [ASM]. Table 2 gives the values for the four classes that cover most of our landscapes, and Section S1 of Online Resource 1 gives the rest and how each source maps onto NLCD classes.

2.6 Host progression: stage durations and clinical mobility

Four mule deer sources put the time from infection to death at 1.5–2.1 yr (Section S1 of Online Resource 1), and we adopt 1.75 yr, [LIT]. None resolve the stages, so we apply the $27 : 36 : 17$ exposed:pre-clinical:clinical split of the box-car model of Almberg et al. (2011), which gives $1/\sigma_1 = 0.591$, $1/\sigma_2 = 0.787$ and $1/\alpha = 0.372$ yr, so

$$\sigma_1 = 1.693 \text{ yr}^{-1}, \qquad \sigma_2 = 1.271 \text{ yr}^{-1}, \qquad \alpha = 2.691 \text{ yr}^{-1}.$$

For clinical mobility, Barrile et al. (2024a) tracked 179 collared mule deer in central Wyoming and found that infected animals began to slow 122 days (110–131) before death, close to clinical onset. Their paper does not give the size of the slowdown as a number, so we estimate it from their data (Barrile et al., 2024b). We divide each infected deer's daily mean speed by that of uninfected deer from the same herd on the same day, which removes seasonal changes in movement. Relative to its level long before death, this ratio falls to about 0.57 in the last ten days of life. Since diffusivity scales with the square of speed at fixed persistence, we average the squared ratio over the last $1/\alpha = 136$ days, which gives 0.68, and adopt $D_R^{\text{scale}} = 0.7$, [LIT]. A scalar multiplier does not capture the changes in habitat selection that happen on similar timescales.

2.7 Prion shedding rates and the prion-unit

Shedding rates cannot be estimated independently of β_2, since scaling ρ_I and ρ_R up and β_2 down by the same constant leaves the dynamics unchanged. Almberg et al. note this, and FitzGerald and Keener (2025) show it for the system of Miller et al. (2006). We fix the scale by requiring each infected animal to deposit one prion-unit (PU)·yr over its shedding career,

$$\frac{\rho_I}{\sigma_2} + \frac{\rho_R}{\alpha} \;=\; 1, \tag{11}$$

[DEF].

The split of deposition between the two stages is set by the carcass. Almberg et al. (2011) have animals shed 2 prion units per week through both shedding phases, with no increase in the living animal, and release a further 100 at death. Since R is the only compartment whose exit is death, ρ_R must carry this pulse. On their 36 : 17 split, a pre-clinical animal deposits 72 units and a clinical one 134. Since our stage durations keep this ratio, $\rho_R = 4\rho_I$ reproduces their $35 : 65$ division of each animal's deposit between the stages, and with (11) gives

$$\rho_I = 0.440, \qquad \rho_R = 1.760 \qquad \text{PU}/(\text{deer} \cdot \text{yr}),$$

[LIT]. The carcass then accounts for 100 of the 206 units an infected deer deposits, and 65% of the total falls in the stage a sight-based cull can see. With the route ratio of Section 2.9, this ratio sets the size of the blind spot, and Section 7.6 tests both. Background mortality would move some of the pulse into the pre-clinical stage and lower the ratio. The model also credits a cull with averting the pulse, which holds only if the carcass is removed from the landscape. A carcass left in the field still releases its pulse, and the cull then averts only the living animal's shedding, about a quarter of the 134 units.

The carcass argument does not apply to direct transmission, which needs a live shedder. We follow Almberg et al. in taking the two stages to be equally infectious, so (6) carries a single β_1 on $I + R$, [LIT]. This is a simplification made for lack of data, since clinical animals salivate excessively but are also emaciated and barely move.

2.8 Environmental persistence δ: the swept parameter

Published values of the environmental decay rate span more than an order of magnitude. Sharp and Pastor (2011) adopt a soil residence of three to ten years in the system of Miller et al. (2006), or 0.10–0.33 yr^{-1}. The half-lives swept by Almberg et al. (2011) give 0.089–2.77, and Potapov et al. (2016) argue that prions become inaccessible after one to two years, giving 0.5–1.0. Garlick et al. (2014) attribute 2.56 to Miller et al. (2006), although the paddocks of Miller et al. (2004) were still infectious two years after the animals were removed. Section S1 of Online Resource 1 sets out the conversions. Given this spread, we sweep

$$\delta \in \{0.125,\ 0.25,\ 0.5,\ 1.0\}\ \text{yr}^{-1},$$

a $\log_2$ ladder of mean persistence from eight years down to one, [LIT] in range. The two lowest rungs bracket the reference value $\delta = 0.15$ at which the route split is set, and since holding R_0 fixed ties β_2 to δ, the sweep varies the generation interval alone (Section 2.9).

2.9 Transmission calibration

Published per-capita transmission coefficients do not transfer to a density-dependent, per-area formulation, so we fix both β through R_0. The split between routes comes from the modeling literature and the conversion to best habitat from the mesh, while the total, the least determined quantity in the model, is swept.

In a wholly susceptible landscape at $S = K_0$, an animal entering E reaches both shedding stages with certainty, so

$$R_0^{\text{dir}} = K_0\,\beta_1 \left(\frac{1}{\sigma_2} + \frac{1}{\alpha}\right), \qquad R_0^{\text{env}} = \frac{K_0\,\beta_2}{\delta}\left(\frac{\rho_I}{\sigma_2} + \frac{\rho_R}{\alpha}\right), \qquad R_0 = R_0^{\text{dir}} + R_0^{\text{env}}, \tag{12}$$

where the bracket in R_0^{env} equals one by (11). We take the ratio $R_0^{\text{env}}/R_0^{\text{dir}}$ from the coefficients of Miller et al. (2006) as rendered by Garlick et al. (2014) and Sharp and Pastor (2011), whose

shedding–transmission products $\gamma_e\epsilon$ agree to within 2% (0.0891 and 0.0874 yr^{-1}), and we use the latter. The direct coefficient is less settled. Miller et al. fit captive deer and scaled infectivity down tenfold for free-ranging animals, giving 0.326 and 0.0326 yr^{-1}, while Garlick et al. need 0.23–0.69 to reproduce the observed La Sal–Abajo spread. We adopt $\beta_{\mathrm{ref}} = 0.1$ yr^{-1}, near the geometric mean of the two renderings, [ASM] in value and [LIT] in range. At $\delta = 0.15$ this gives $\gamma_e\epsilon/(\delta\beta_{\mathrm{ref}}) = 5.82$, an environmental share of 85%, which we hold fixed across the sweep.

The split is sensitive to β_{ref}. The same formula gives 18.2 at the free-ranging rate of Miller et al., 2.6 at the rate Garlick et al. calibrate to observed spread, and 1.8 at the captive rate. Still, environmental dominance is not specific to these renderings, since Yin et al. (2026), using an independent parameter set, find $R_0^{\mathrm{env}} \gg R_0^{\mathrm{dir}}$ across most of their parameter plane. Weighting each route by intensity times residence puts 60% of R_0 in the clinical class and 40% in the pre-clinical one, a division that is central to the control problem (Section 3).

Since R_0 is linear in $K = K_0 c(\mathbf{x})$, it is a local quantity here. Writing

$$R_0(\mathbf{x}) = \widetilde{R_0}\, c(\mathbf{x}), \qquad \overline{c} = \frac{1}{|\Gamma|}\int_\Gamma c(\mathbf{x})\,\mathrm{d}\Gamma, \tag{13}$$

$\widetilde{R_0}$ is the value in the best cover class and $\widetilde{R_0}\,\overline{c}$ the area average. On the three meshes of Section 7.5, $\overline{c} = 0.78$–0.84. Here $\widetilde{R_0}$ is a *calibration index*, not the invasion threshold of (1)–(5), which is the sign of the principal eigenvalue of the linearization about the disease-free state and depends on the diffusion tensor and the spatial profile of K. The index is also evaluated at $S = K(\mathbf{x})$ rather than at the disease-free equilibrium $S_\infty(\mathbf{x})$, the steady state of (1) with $E = I = R = W = 0$. Diffusion smooths S_∞ across habitat boundaries, so on a fragmented mosaic the index misstates local transmission near those boundaries. We use $\widetilde{R_0}$ only to fix β_1 and β_2 and to index the sweep. The distinction from $\widetilde{R_0}\,\overline{c}$ still matters, since reading a landscape average as the best-habitat value puts the index below one over much of a poor mosaic, where the epizootic then fails to establish.

Inverting (12) at the adopted stage rates, with $K_0 = 3$ and the 5.82:1 split, gives

$$\beta_1 \;=\; 0.0422\,\widetilde{R_0}, \qquad \beta_2 \;=\; 0.2845\,\widetilde{R_0}\,\delta. \tag{14}$$

Since R_0^{env} carries a factor of $1/\delta$, holding R_0 fixed makes β_2 proportional to the decay rate. This coupling is deliberate. It holds R_0, the route split and the host endemic equilibrium fixed across the sweep, so lifetime transmission per infected animal is the same and only its *timing* differs. The same R_0 then produces a faster epizootic at larger δ. The design suppresses an effect that is probably real (Almberg et al. (2011) find R_0 rising by about 0.29 per additional year of prion half-life) in order to isolate the timing effect (Section S2 of Online Resource 1).

The uniform endemic state in the best cover class is also a function of $\widetilde{R_0}$ and δ. Writing

$$F(\widetilde{R_0}) \;=\; \frac{0.78\,\bigl(1 - 1/\widetilde{R_0}\bigr)}{\widetilde{R_0} + 0.455}$$

for the equilibrium infection rate in deer/(km^2·yr), the endemic state is

$$S^* = \frac{3}{\widetilde{R_0}}, \qquad E^* = 0.591\,F, \qquad I^* = 0.787\,F, \qquad R^* = 0.372\,F, \qquad W^* = \frac{F}{\delta}, \tag{15}$$

where each infected stage holds its inflow for its mean duration and the reservoir balances deposition against decay. Only W^* depends on δ, and by (11) it does not depend on ρ_R/ρ_I. At the low anchor $\widetilde{R_0} = 2$, $S^* = 1.5$ deer/km^2.

Field evidence gives a range for R_0. Inverting (15) against three readings from Geremia et al. (2015) gives $R_0 = 1.15$–1.34, a lower bound since that herd is hunted below capacity, and the roughly one adult in four infected at Table Mesa (Miller et al., 2008) implies $R_0 \approx 2.9$, [CAL]. This range is probably low, because a persistent reservoir makes slow prevalence growth compatible with a large R_0. When constrained to growth rates observed in Colorado, Wyoming and Wisconsin, Almberg et al. (2011) obtain R_0 as high as 10–11. We therefore sweep R_0 over 2, 3.5 and 5, taking 2 as the low reference and extending toward the range of Almberg et al. (Section 7.1). Since a larger R_0 makes the same effort buy less, the low anchor makes the recoverable fractions of Section 7 optimistic. Section S1 of Online Resource 1 gives the inversions and a structural prevalence ceiling of 31%.

2.10 Stability of the endemic state

Sharp and Pastor (2011) showed that a model with this structure undergoes a Hopf bifurcation as carrying capacity rises, producing recurrent epizootics every ten to fifteen years. In the full model, the endemic equilibrium in the best cover class is stable at 11 of the 12 grid cells, with decay e-folding times of 14 yr or longer, and on all three landscapes for every carrying capacity the land-cover map assigns. At ($\delta = 1$, $\widetilde{R_0} = 5$) a mode of period 22 yr grows with an e-folding time of about 350 yr, which is negligible over the 30 yr horizon.

Table 1 Model parameters, adopted values and provenance labels (Section 2.4). Shaded rows are parameters that the literature does not determine for this region, and are either assumed or swept. Rows marked *swept* are varied across the runs of Section 7, and both transmission coefficients move with $\widetilde{R_0}$ by (14). Objective and discretization parameters are in Section 3.3.

Symbol	Description	Value	Units	Provenance & Section
σ_1	Exposed → pre-clinical rate	1.693	yr^{-1}	[LIT] 2.6
σ_2	Pre-clinical → clinical rate	1.271	yr^{-1}	[LIT] 2.6
α	CWD removal rate of clinical deer	2.691	yr^{-1}	[LIT] 2.6
δ	Environmental prion decay	*swept* $0.125, 0.25, 0.5, 1.0$	yr^{-1}	[LIT] 2.8
ρ_I	Pre-clinical shedding rate	0.440	PU/(deer·yr)	[DEF] 2.7, (11)
ρ_R	Clinical shedding rate	1.760	PU/(deer·yr)	[LIT] 2.7, (11)
$\widetilde{R_0}$	Best-habitat reproductive number	*swept* $2, 3.5, 5$	—	[CAL] 2.9, (13)
β_1	Direct transmission, both stages	$0.0422\,\widetilde{R_0}$	km^2/(deer·yr)	[CAL] 2.9, (12), (14)
β_2	Environmental transmission	$0.2845\,\widetilde{R_0}\,\delta$	km^2/(PU·yr)	[CAL] 2.9, (12), (14)
r	Intrinsic growth rate	0.26	yr^{-1}	[LIT] 2.4
K_0	Base carrying capacity	3	deer/km^2	[ASM] 2.4
κ	Diffusivity scale	4×10^6	m^2/yr	[ASM] 2.5
$D^{\mathrm{scale}}_{S,E,I,R}$	Compartment mobilities	1, 1, 1, 0.7	—	[LIT] 2.6
γ	Slope penalty rate	4.7	—	[LIT] 2.5
$c(\mathbf{x})$, $p(\mathbf{x})$	Land-cover coefficients	Table 2	—	mixed; 2.5
ι	Tensor degeneracy floor	10^{-3}	—	[NUM] 2.5

3 An Optimal Control Problem: Targeted Culling

We now use the model as the constraint in an optimal control problem: given a planning horizon and a culling effort that can vary over the landscape and in time, where and when should that effort be spent?

Table 2 The four land-cover classes that cover most of the landscape meshes of Section 7.5, with the share of each mesh they cover. Column c is the habitat quality in (8), from the mule deer motilities of Garlick et al. (2014), and p the permeability in (9), from the log-diffusion coefficients of Mandujano Reyes et al. (2025). All eight values are [LIT]. Shares are of the NLCD surface area of each 100 km box. The sweep disk uses $c = p = 1$. Section S1 of Online Resource 1 gives the other classes and the readings behind every value.

				Share of mesh (%)		
Code	Land cover class	c	p	Upper Arkansas	Carbondale	Estes Park
42	Evergreen forest	0.85	0.62	40.2	34.5	35.0
52	Shrub or scrub	1.00	0.81	29.6	29.8	21.3
71	Grassland or herbaceous	0.80	0.80	16.3	1.2	25.6
41	Deciduous forest	0.75	1.00	2.9	25.2	0.4
	Other classes			11.0	9.3	17.7

3.1 The controlled system and the objective

The control $v : \Gamma \times (0,T) \to \mathbb{R}$ is a per-capita removal rate (yr^{-1}) applied to the clinical class, so (4) becomes

$$\dot{R} = \nabla_\Gamma \cdot \big(\mathbf{D}_R \nabla_\Gamma R\big) + \sigma_2 I - \alpha R \;-\; v\,R. \tag{16}$$

This represents a sight-based cull which is blind to the exposed and pre-clinical classes. We call (1)–(5), with (4) replaced by (16), the no-flux conditions (7) and initial data y^0, the *controlled system*.

We assume that control effort is bounded by some $v_{\max} > 0$, so that controls live in the set

$$\mathcal{V} = \Big\{ v \in L^\infty\big(\Gamma \times (0,T)\big) \;:\; 0 \le v \le v_{\max} \;\text{ a.e.} \Big\}. \tag{17}$$

Since no real culling effort changes continuously, we optimize over controls that are constant on each of N_b blocks $[t_{b-1}, t_b)$ partitioning a *control horizon* $[0, T_\mathrm{c}]$. This restriction also lets the quasi-Newton method of Section 6.8 fit in memory. The admissible set is then

$$\mathcal{V}^\mathrm{b} = \Big\{ v \in \mathcal{V} \;:\; v(\mathbf{x},t) = v_b(\mathbf{x}) \;\text{ for } t \in [t_{b-1}, t_b),\; b = 1, \dots, N_\mathrm{b}; \quad v \equiv 0 \text{ on } (T_\mathrm{c}, T] \Big\}. \tag{18}$$

The uncontrolled tail $(T_\mathrm{c}, T]$ takes the place of a terminal cost. Without this tail, infections generated near the end of the *control horizon* would not be charged in full, allowing controls to stop effort early at little cost.

The objective functional is

$$J(y,v) \;=\; \int_0^T \!\!\int_\Gamma \Big[\, c_1\big(S_\infty - S\big) \;+\; c_2\, W \;+\; c_3\, v^2 \;+\; c_4\, v \,\Big]\, \mathrm{d}\Gamma\, \mathrm{d}t \tag{19}$$

with $c_1, c_2, c_4 \ge 0$ and $c_3 > 0$, where $S_\infty(\mathbf{x})$ is the disease-free equilibrium. The terms are the accumulated *susceptible deficit*, the *accumulated prion load* and an *elastic net* penalty on effort. The two state terms cover the two natural ways of pricing the epizootic: by its damage to the herd and by the infection itself. The deficit counts infection only through the susceptibles it removes. The prion load counts the shedding population directly. Since W neither diffuses nor is controlled, (5) integrates exactly to $\int_0^T W\,\mathrm{d}t = \delta^{-1}\int_0^T (\rho_I I + \rho_R R)\,\mathrm{d}t - \delta^{-1}\big(W(T) - W(0)\big)$, so pricing W is pricing $\rho_I I + \rho_R R$ up to the constant $W(0)/\delta$ and a credit for the contamination left at T. The

two terms also steer the control differently, since the deficit reaches the clinical adjoint only where susceptibles remain (Section 4.2).

Writing $y(v)$ for the solution of the controlled system and $\hat{J}(v) := J(y(v), v)$ for the reduced objective, the problem is

$$\min_{v \in \mathcal{V}^{\mathrm{b}}} \ \hat{J}(v). \tag{20}$$

A minimizer exists by the argument of Miller Neilan and Lenhart (2011, Thm. A.2) for a parabolic susceptible–infected–recovered (SIR) system of the same type, since $\mathcal{V}^{\mathrm{b}}$ is bounded, closed and convex and $c_3 > 0$ makes the control cost weakly lower semicontinuous. It need not be unique. Since $v \mapsto y(v)$ is nonlinear, $\hat{J}$ is nonconvex, and the method of Section 6.8 finds a stationary point.

3.2 The reach of the control

Restricting the control to the clinical class is central to this paper. At any instant the clinical class holds under a third of the shedding animals, and every clinical animal has already spent about forty weeks transmitting as a pre-clinical one. The control also cannot act on W directly, so contamination can be prevented but never removed.

We can see what such a control can achieve in closed form. A constant cull enters only through $\alpha \to \alpha + v$, so in (12) it shortens the clinical terms and leaves the two $1/\sigma_2$ terms, the transmission from the pre-clinical stage, unchanged. Splitting $R_0 = R_0^{\mathrm{pre}} + R_0^{\mathrm{clin}}$ by stage, the control reproduction number is

$$R_c(v) \ = \ R_0^{\mathrm{pre}} \ + \ R_0^{\mathrm{clin}} \, \frac{\alpha}{\alpha + v}, \tag{21}$$

which decreases toward R_0^{pre} as v grows but never reaches it. A cull of the clinical class can therefore bring R_c below one only where $R_0^{\mathrm{pre}} < 1$ and

$$v_{\max} \ > \ \alpha \, \frac{R_0 - 1}{1 - R_0^{\mathrm{pre}}}. \tag{22}$$

The pre-clinical fraction depends on the route ratio $m = R_0^{\mathrm{env}} / R_0^{\mathrm{dir}}$, the deposition ratio $\varrho = \rho_R / \rho_I$ and the duration ratio $\tau = \sigma_2 / \alpha$. Within the direct route the stages divide in proportion to their durations, and within the environmental route in proportion to their deposits, so

$$\frac{R_0^{\mathrm{pre}}}{R_0} \ = \ \frac{1}{1+m} \left[\frac{1}{1+\tau} \ + \ \frac{m}{1+\varrho\,\tau} \right]. \tag{23}$$

At our adopted values of $m = 5.82$, $\varrho = 4$ and $\tau = 0.472$ this fraction is 0.395, so $R_0^{\mathrm{pre}} < 1$ requires $R_0 < 2.53$. At $R_0 = 2$ the condition requires $v_{\max} > 12.8$ yr^{-1}, a cull that takes 83% of clinical animals before they die and is likely well beyond what a sight-based protocol can sustain (Section 3.3), and at $R_0 = 3.5$ and 5 no cull rate suffices. Figure 1 maps this condition over R_0 and ϱ. Raising ϱ from 4 to 8 halves the rate needed at $R_0 = 2$, to 6.1 yr^{-1}, but at our cap of $v_{\max} = 3$ yr^{-1} and the adopted m, the index stays above one at every cell of the sweep whatever the value of ϱ.

Equation (23) does not depend on the spatial model. It shows how the ratios we are least sure of set the local reach of a sight-based cull. Raising ϱ, or raising m while $\varrho > 1$, shifts transmission toward the clinical stage and lowers R_0^{pre}/R_0, so the same cull reaches more of R_0. Taking $\varrho \geq 1$ is safe, since clinical animals shed at least as much as pre-clinical ones while alive and ρ_R also carries the carcass pulse (Section 2.7). Each ratio could in principle be estimated in the field, τ from the timing of clinical onset in collared animals, ϱ from shedding and carcass prion loads, and m, the hardest, from separate estimates of β_1 and β_2. With an estimate of R_0, (23) and (22) then

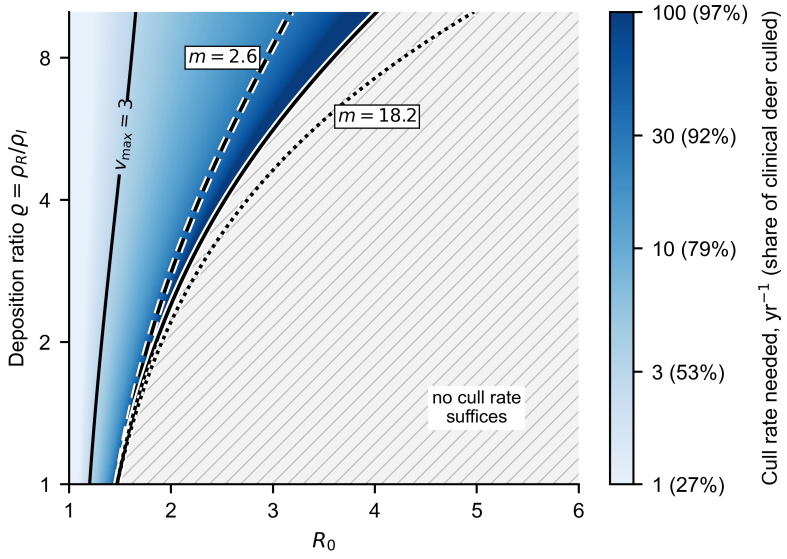


Fig. 1 The smallest constant cull rate that brings the control reproduction number (21) below one, from (22), over R_0 and the deposition ratio ϱ at the adopted route ratio $m = 5.82$. The color scale also gives the share of clinical deer culled before they die, $v/(v + \alpha)$, and the bold line marks the cap $v_{\max} = 3$. In the hatched region $R_0^{\text{pre}} \geq 1$ and no cull rate suffices. Its edge is solid at $m = 5.82$, dashed at $m = 2.6$ and dotted at $m = 18.2$, the route ratios tested in Section 7.6

test whether a sight-based cull could stop an outbreak before any spatial model is built. Like $\widetilde{R_0}$, however, (21) is a local well-mixed index rather than the spatial threshold. It explains why every result in Section 7 is a reduction in accumulated burden and never an elimination, and what the spatial optimization adds is where and when a bounded effort does the most good.

3.3 Control problem parameters

None of the weights is an economic cost. Setting c_1:c_3:c_4 would require pricing a culled deer against an infected one in a shared currency, and neither price exists. Uehlinger et al. (2016) found no study that supplies a per-animal cost of CWD management. The ratios instead state what is being traded off. They support conclusions about the spatial and temporal structure of the optimal policy and comparisons under identical weights, but not about absolute effort levels or cost-effectiveness.

We set $c_2 = 0$ except in one comparison (Section 7.3). Then c_1 only sets the overall scale of J, which does not change the minimizer, so we fix $c_1 = 10$ and read c_1/c_3 and c_1/c_4 directly off c_3 and c_4.

The quadratic weight c_3 is a Tikhonov term. It makes J strictly convex in the control at fixed state and keeps the optimal control a continuous function of the adjoint rather than bang-bang, although $\hat{J}$ stays nonconvex. Since only c_1/c_3 matters, c_3 is also the price of effort, so we calibrate it rather than choose it. The value $c_3 = 0.5$ is set against the capacity $v_{\max}$ by the saturation probe of Section S2 of Online Resource 1, [CAL].

The linear weight c_4 is an L^1 term, [ASM], that makes the optimal effort exactly zero wherever the benefit density falls below it (Section 5). At $c_4 = 0$ effort is positive wherever clinical animals remain, so a purely quadratic objective returns a thin cull everywhere, which no agency can carry out. The L^1 term instead separates zones that are culled from zones that are not. We fix $c_4 = 0.5$. The same policy follows from a budget on the total effort over the control window, whose multiplier

takes the place of c_4 in the optimality condition. The c_4 form prescribes the threshold and finds the effort, while the budget prescribes the effort and finds the threshold. We use the first since it needs no outer iteration on the multiplier, although the second is likely closer to a real culling constraint.

We adopt $v_{\max} = 3\ \mathrm{yr}^{-1}$, [ASM], as our estimate of what a sustained cull at full effort could achieve. Against $\alpha = 2.69\ \mathrm{yr}^{-1}$, this cap culls $v/(v+\alpha) = 53\%$ of clinical animals before they die of the disease and shortens the mean clinical stage from 136 to 64 days. Clinical animals are conspicuous and barely move, so they neither hide nor leave. The carcass pulse enters $\rho_R R$ as the pulse size times the natural death flux αR, so animals leaving at the cull rate vR deposit no pulse. A clinical animal therefore deposits $\rho_R/(\alpha+v)$, its live shedding over the shortened stay plus a pulse with probability $\alpha/(\alpha+v)$, and we read v as a cull that removes the carcass.

We take T_{tail} to be one infect–deposit–decay cycle at the slowest grid cell, $1/\sigma_1 + 1/\sigma_2 + 1/\alpha = 1.749$ yr plus $1/\delta = 8$ yr at $\delta = 0.125$, so $T_{\mathrm{tail}} = 10$ yr, [DEF]. We use the same T_{tail} at every cell so that objective values stay comparable. The tail makes the terminal term depend on δ through the reservoir transient, and it still truncates the future. Section 7.6 tests the effect of doubling it. Table 3 collects these settings with the horizon and the discretization.

Table 3 Objective, horizon and discretization parameters. Shaded rows are assumed. Every run uses these values except the comparison of Section 7.3, which moves the state weight from c_1 onto c_2.

Symbol	Description	Value	Units	Provenance & Section
c_1	Weight on the susceptible deficit	10	—	[DEF] 3.3
c_2	Weight on the environmental load	0	—	[DEF] 7.1
c_3	Quadratic weight on v	0.5	—	[CAL] 3.3
c_4	Sparsity (L^1) weight on v	0.5	—	[ASM] 3.3
$v_{\max}$	Maximum effort on R	3	yr^{-1}	[ASM] 3.3
T_{c}	Control horizon	20	yr	[ASM] 7.1
T_{tail}	Continuation tail	10	yr	[DEF] 3.3
T	Simulated horizon	30	yr	[DEF] 3.1
N_{b}	Control blocks	20	—	[ASM] 7.1
Δt	Time step	0.05	yr	[NUM] 6.3

4 Derivation of the Adjoint System

To optimize (20) with a gradient-based method we need $\hat{J}'(v)$. We begin by collecting the state as $y = (y_S, y_E, y_I, y_R, y_W) := (S, E, I, R, W)$, so that the reaction terms, including the control, are

$$\begin{aligned} f_S(y,v) &= G(y) - S\Lambda, & G(y) &:= rS\Big(1 - \frac{S+E+I+R}{K}\Big), \\ f_E(y,v) &= S\Lambda - \sigma_1 E, & f_I(y,v) &= \sigma_1 E - \sigma_2 I, \\ f_R(y,v) &= \sigma_2 I - (\alpha+v)R, & f_W(y,v) &= \rho_I I + \rho_R R - \delta W. \end{aligned} \tag{24}$$

With $\mathbf{D}_W := \mathbf{0}$, the controlled system reads

$$\dot{y}_k = \nabla_\Gamma \cdot \big(\mathbf{D}_k \nabla_\Gamma y_k\big) + f_k(y,v), \qquad k \in \{S,E,I,R,W\}. \tag{25}$$

The only nonzero control derivative is $\partial f_R/\partial v = -R$, in the one compartment the cull can see.

4.1 The Lagrangian and its derivative

Adjoining the state equations with multipliers $\lambda = (\lambda_S, \lambda_E, \lambda_I, \lambda_R, \lambda_W)$, we form the Lagrangian

$$\mathcal{L}(y, v, \lambda) \;=\; J(y, v) \;-\; \sum_k \int_0^T \!\! \int_\Gamma \lambda_k \Big[\dot{y}_k - \nabla_\Gamma \cdot \big(\mathbf{D}_k \nabla_\Gamma y_k \big) - f_k(y, v) \Big] \, \mathrm{d}\Gamma \, \mathrm{d}t. \tag{26}$$

For any y that satisfies the state equations, $\mathcal{L} = J$ for every λ, so we are free to choose λ to eliminate the state sensitivities. With the running cost and Hamiltonian density

$$\ell(y, v) := c_1 \big(S_\infty - S \big) + c_2 W + c_3 v^2 + c_4 v, \qquad \mathcal{H}(y, v, \lambda) := \ell(y, v) + \sum_k \lambda_k f_k(y, v), \tag{27}$$

we perturb $v \to v + \varepsilon \, \delta v$ within $\mathcal{V}$ and let δy solve the linearized system

$$\delta \dot{y}_k - \nabla_\Gamma \cdot \big(\mathbf{D}_k \nabla_\Gamma \delta y_k \big) - \sum_j \frac{\partial f_k}{\partial y_j} \, \delta y_j \;=\; \frac{\partial f_k}{\partial v} \, \delta v, \tag{28}$$

with $\delta y_k(\cdot, 0) = 0$ and no flux. Taking the Gâteaux derivative of (26), substituting (28) and integrating by parts once in time and twice in space (Hinze et al., 2009, Ch. 1), with $\lambda_k(\cdot, T) = 0$ and $(\mathbf{D}_k \nabla_\Gamma \lambda_k) \cdot \boldsymbol{\nu} = 0$ to remove the boundary terms, leaves

$$\delta \mathcal{L} = \sum_k \int_0^T \!\! \int_\Gamma \delta y_k \Big[\dot{\lambda}_k + \nabla_\Gamma \cdot \big(\mathbf{D}_k \nabla_\Gamma \lambda_k \big) + \frac{\partial \mathcal{H}}{\partial y_k} \Big] \, \mathrm{d}\Gamma \, \mathrm{d}t \;+\; \int_0^T \!\! \int_\Gamma \frac{\partial \mathcal{H}}{\partial v} \, \delta v \, \mathrm{d}\Gamma \, \mathrm{d}t. \tag{29}$$

Choosing λ so that the first bracket vanishes leaves a derivative in δv alone.

4.2 The adjoint system

With $\xi := r(1 - N/K) - rS/K$ and $\omega := -rS/K$, the adjoint system is

$$-\dot{\lambda}_S - \nabla_\Gamma \cdot \big(\mathbf{D}_S \nabla_\Gamma \lambda_S \big) = -c_1 + \lambda_S \big[\xi - \Lambda \big] + \lambda_E \, \Lambda, \tag{30}$$

$$-\dot{\lambda}_E - \nabla_\Gamma \cdot \big(\mathbf{D}_E \nabla_\Gamma \lambda_E \big) = \lambda_S \, \omega - \lambda_E \, \sigma_1 + \lambda_I \, \sigma_1, \tag{31}$$

$$-\dot{\lambda}_I - \nabla_\Gamma \cdot \big(\mathbf{D}_I \nabla_\Gamma \lambda_I \big) = \lambda_S \big[\omega - \beta_1 S \big] + \lambda_E \, \beta_1 S - \lambda_I \, \sigma_2 + \lambda_R \, \sigma_2 + \lambda_W \, \rho_I, \tag{32}$$

$$-\dot{\lambda}_R - \nabla_\Gamma \cdot \big(\mathbf{D}_R \nabla_\Gamma \lambda_R \big) = \lambda_S \big[\omega - \beta_1 S \big] + \lambda_E \, \beta_1 S - \lambda_R \big(\alpha + v \big) + \lambda_W \, \rho_R, \tag{33}$$

$$-\dot{\lambda}_W = c_2 - \lambda_S \, \beta_2 S + \lambda_E \, \beta_2 S - \lambda_W \, \delta, \tag{34}$$

with no-flux boundary conditions and $\lambda_k(\cdot, T) = 0$. Over non-habitat cells ($K < 0.01$) the logistic term of (1) is replaced by $-rS$, giving $\xi = -r$ and $\omega = 0$.

The adjoint $\lambda_k(\mathbf{x}, t)$ is the marginal cost over $[t, T]$ of one additional unit of compartment k. For times in the control window, the part of $[t, T]$ beyond T_c is the continuation tail, whose cost reaches the last controlled block through $\lambda(\cdot, T_\mathrm{c})$.

The running costs enter (30)–(34) only as source terms, $-c_1$ in row S and c_2 in row W. Given the trajectory, the adjoint is linear, so it is the sum of the responses to these two sources, and they reach λ_R by different routes. Every coupling out of λ_S carries a factor of S (the entries feeding rows E, I, R and W are ω, $\omega - \beta_1 S$, $\omega - \beta_1 S$ and $-\beta_2 S$), because S multiplies every cross-compartment term of (1). The couplings out of λ_W carry no such factor, feeding rows I and R through the shedding rates ρ_I and ρ_R. The deficit therefore passes into the infected adjoints only where susceptibles remain, while a reservoir weight would pass into λ_R wherever clinical animals

shed, the depleted core included. Neither route makes λ_R proportional to S or R pointwise, since each source then travels backward through a coupled parabolic system. The optimal effort is $v^* = P[(\lambda_R R - c_4)/(2c_3)]$ (Section 5), so all of the steering lies in λ_R, and effort is positive only where $\lambda_R > c_4/R$. With $c_2 = 0$ only the S-gated route is active. Section 7.3 tests whether this gating decides where effort falls and shows the core effort that the reservoir route adds.

The weak form of the adjoint follows Section 6.1, with (34) advanced at each node, and its bilinear parts are those of the state problem.

5 Optimality Condition and the Reduced Gradient

With λ solving (30)–(34), only the Hamiltonian derivative survives in (29), which gives the reduced gradient

$$\nabla \hat{J}(v) \;=\; 2c_3\, v + c_4 - \lambda_R\, R, \tag{35}$$

as the Riesz representative in $L^2(\Gamma \times (0, T_c))$. One state solve and one adjoint solve give the gradient with respect to every control degree of freedom.

The optimizer raises v wherever $\lambda_R R$, the clinical density times the marginal future cost of one clinical animal, exceeds the marginal cost $c_4 + 2c_3 v$ of effort. Effort is worthless where there is nothing visible to cull ($R \to 0$) or where culling buys nothing ($\lambda_R \to 0$), and the continuation tail keeps the latter from happening at the end of the control window. Notably, neither $\lambda_I I$ nor $\lambda_W W$ appears. The optimizer reaches the pre-clinical class and the reservoir only indirectly, by removing the clinical animals that I feeds and that contaminate W, and Section 7 measures what that indirect effect is worth.

5.1 First-order necessary condition

Since $\mathcal{V}$ is convex, a local minimizer v^* satisfies

$$\langle \nabla \hat{J}(v^*), \tilde{v} - v^* \rangle \;\geq\; 0 \qquad \text{for all } \tilde{v} \in \mathcal{V}, \tag{36}$$

which for a pointwise box is equivalent to

$$v^*(\mathbf{x}, t) = P_{[0, v_{\max}]}\!\left(\frac{\lambda_R R - c_4}{2c_3} \right), \tag{37}$$

where y and λ are those of v^*, so the condition is coupled through the state and adjoint solves. Over the block-constant set $\mathcal{V}^{\mathrm{b}}$, testing (36) against block-constant variations replaces the instantaneous benefit density by its block average,

$$v_b^*(\mathbf{x}) \;=\; P_{[0, v_{\max}]}\!\left(\frac{\overline{\chi}_b(\mathbf{x}) - c_4}{2c_3} \right), \qquad \overline{\chi}_b(\mathbf{x}) \;:=\; \frac{1}{\Delta t_b} \int_{t_{b-1}}^{t_b} \lambda_R(\mathbf{x}, t)\, R(\mathbf{x}, t)\, \mathrm{d}t, \tag{38}$$

with $\Delta t_b = t_b - t_{b-1}$ the block length, which is the continuous counterpart of the block gradient in Section 6.5. Every statement about the runs refers to (38), while (37) serves for intuition. Pointwise, c_4 acts as a cutoff on the benefit density

$$\chi(\mathbf{x}, t) \;:=\; \lambda_R\, R. \tag{39}$$

The optimal effort is exactly zero wherever $\chi < c_4$ and above that ramps up with slope $1/(2c_3)$, so c_4 shifts χ and c_3 sets the slope, which is why the two can be set independently.

Any zero set has to come from c_4. The sign structure of (30)–(34) suggests that $\lambda_R > 0$ on $[0, T)$. We have $\lambda_S < 0$, and since $\omega \le 0$, the forcing $\lambda_S[\omega - \beta_1 S] + \lambda_E \beta_1 S$ of row R is nonnegative when $\lambda_E \ge 0$. Row W has forcing $(\lambda_E - \lambda_S)\beta_2 S \ge 0$, so $\lambda_W \ge 0$, and a backward parabolic equation with nonnegative forcing and zero terminal data has a nonnegative solution. This argument is formal, since the signs of λ_S and λ_E are read off the same coupled system. A proof would need a maximum principle for the cooperative transposed system, which we do not give. Taken at face value, it says that a sight-based cull is never harmful, only sometimes not worth its cost, so at $c_4 = 0$ there is no zero set. This is consistent with the disk runs, where the zero set at $c_4 = 0.5$ lies ahead of the invasion, where $R \approx 0$, and effort in the depleted core never falls below 0.13 yr^{-1} (Section 7.3).

In the code we monitor the equivalent fixed-point form

$$v^* = P_{\mathcal{V}}\big(v^* - s\,\nabla \hat{J}(v^*)\big) \quad \text{for any } s > 0, \tag{40}$$

with the residual measured at $s = 1$ in the L^2 norm. Condition (37) suggests the forward–backward sweep of Lenhart and Workman (2007), but that method has no descent guarantee and can diverge when the coupling is strong, as it is here through the $-\lambda_R(\alpha + v)$ term. The method of Section 6.8 keeps the forward–backward structure but takes a safeguarded descent step.

6 Numerical Methods and Implementation

The state system, the adjoint system and the gradient (35) share one surface finite element space and one lumped-mass implicit–explicit (IMEX) step, and since the diffusion tensor is symmetric, the forward and backward sweeps also share one matrix factorization. We build Γ_h by smoothing the 30 m DEM of Section 2.1 with a Gaussian of standard deviation 90 m, subsampling it to a 90 m grid, triangulating, and remeshing the surface with `mmg` under an anisotropic metric built from the slope factor (10). The target edge length is $\sqrt{2 - \eta(\theta)}\, h_{\max}$ along $\mathbf{e}_{\mathrm{c}}$ and $\sqrt{\eta(\theta)}\, h_{\max}$ along $\mathbf{e}_{\mathrm{s}}$, following the diffusion lengths $\sqrt{\kappa(2-\eta)/r}$ and $\sqrt{\kappa\eta/r}$, so elements are long along contours and thin up steep slopes, down to $h_{\min} = \sqrt{\iota}\, h_{\max}$. We take $h_{\max} = 250$ m on all three landscapes, the uniform h of the sweep disk (Section 7.1), giving 323,798–370,354 nodes (Section S3 of Online Resource 1). Each vertex carries the land-cover class of the NLCD cell it falls in, and K and p are the piecewise-linear interpolants of their values at the vertices. A class boundary is therefore smoothed over one element, at most a few hundred meters, well below the flat-ground diffusion length $\sqrt{\kappa/r} = 3.92$ km. The tensor (9) uses the normal of each surface element, so η and the frame $(\mathbf{e}_{\mathrm{s}}, \mathbf{e}_{\mathrm{c}})$ are constant on each element, while the mesh metric uses the area-weighted average of the triangle normals around each vertex. The solver is implemented in FEniCSx with `PETSc`. The authors used Claude (Anthropic) to assist with developing the codebase and with revising and editing drafts of the manuscript. The authors checked all code, results and text and take full responsibility for the content of the article.

6.1 Weak formulation

Let $V_h \subset H^1(\Gamma_h)$ be the continuous piecewise-linear space with nodal basis $\{\varphi_i\}_{i=1}^{N}$. Multiplying the host equations by $\psi \in V_h$ and applying the surface divergence theorem gives the weak problem: find $X_h \in V_h$ such that

$$\int_{\Gamma_h} \dot{X}_h\, \psi \,\mathrm{d}A = -\int_{\Gamma_h} \mathbf{D}_X \nabla_\Gamma X_h \cdot \nabla_\Gamma \psi \,\mathrm{d}A + \int_{\Gamma_h} f_X(y_h, v)\, \psi \,\mathrm{d}A \qquad \forall\, \psi \in V_h, \quad X \in \{S, E, I, R\}, \tag{41}$$

with f_X from (24), and the no-flux conditions are natural. We build Λ once and share it between the S and E rows, so the discrete system conserves infections by construction. Equation (5) is discretized at the nodes.

6.2 Semi-discrete system

Expanding in the basis gives

$$\mathbf{M}\dot{\mathbf{X}} = -\mathbf{A}_X\,\mathbf{X} + \mathbf{F}_X(\mathbf{y}), \quad X \in \{S, E, I, R\}, \qquad \dot{\mathbf{W}} = \rho_I\,\mathbf{I} + \rho_R\,\mathbf{R} - \delta\,\mathbf{W}, \tag{42}$$

with consistent mass matrix $[\mathbf{M}]_{ij} = \int_{\Gamma_h} \varphi_i\varphi_j \,\mathrm{d}A$, stiffness matrix $[\mathbf{A}_X]_{ij} = \int_{\Gamma_h} \mathbf{D}_X\,\nabla_\Gamma\varphi_i\cdot\nabla_\Gamma\varphi_j \,\mathrm{d}A$ built with (9), and $\mathbf{F}_X$ the projected reaction terms.

We replace $\mathbf{M}$ by the lumped mass matrix $\mathbf{M}_\mathrm{L} = \mathrm{diag}(\mathbf{M1})$, obtained with a vertex quadrature rule, which preserves total mass and costs no order of accuracy for linear elements. Lumping removes one source of undershoot, since the positive off-diagonal entries of $\mathbf{M}$ can let the diffusion step go negative on sharp data. It does not guarantee positivity, since for an anisotropic tensor that also requires a mesh that is nonobtuse in the metric of $\mathbf{D}_X$, which our mesh generator does not enforce on steep ground. We therefore claim no discrete maximum principle, although in uncontrolled forward runs on two of our meshes no compartment went negative at any time step. Reaction terms use a Gauss rule.

6.3 Time integration

Diffusion is linear and potentially stiff, while reaction is nonlinear and non-stiff, so we advance each host compartment by a first-order IMEX step,

$$\left(\mathbf{M}_\mathrm{L} + \Delta t\,\mathbf{A}_X\right)\mathbf{X}^{n+1} = \mathbf{M}_\mathrm{L}\,\mathbf{X}^n + \Delta t\,\mathbf{F}_X(\mathbf{y}^n), \qquad X \in \{S, E, I, R\}, \tag{43}$$

and W by the exact solution of its linear ODE at each node,

$$W_i^{n+1} = W_i^n \mathrm{e}^{-\delta\Delta t} + \left(\rho_I I_i^n + \rho_R R_i^n\right)\frac{1 - \mathrm{e}^{-\delta\Delta t}}{\delta}, \tag{44}$$

which is unconditionally stable. The time step is limited by accuracy rather than stability, since the fastest explicit rate gives $\alpha\Delta t = 0.135$ at $\Delta t = 0.05$ yr, well below the explicit-Euler stability limit of 2. We check convergence in Δt in Section 7.6.

6.4 Monolithic assembly and block factorization

All five compartments share a mixed space of degree-1 Lagrange elements. The implicit side of (43) is the bilinear form

$$a(Q, \Psi) = \sum_{X\in\{S,E,I,R\}} \int_{\Gamma_h} \Big[\, q_X\psi_X + \Delta t\,(\mathbf{D}_X\,\nabla_\Gamma q_X)\cdot\nabla_\Gamma\psi_X \Big]\,\mathrm{d}A \;+\; \int_{\Gamma_h} q_W\psi_W\,\mathrm{d}A, \tag{45}$$

where Q and Ψ are the mixed trial and test functions. Its coefficients do not depend on the state, so we assemble and factorize its matrix $\boldsymbol{\mathcal{A}}$ once before the time loop and rebuild only the load vector. Since there are no cross-compartment terms, $\boldsymbol{\mathcal{A}} = \mathrm{blockdiag}(\mathbf{M}_\mathrm{L} + \Delta t\mathbf{A}_S,\ \ldots,\ \mathbf{M}_\mathrm{L} + \Delta t\mathbf{A}_R,\ \mathbf{M}_\mathrm{L})$ is factorized block by block, and at the adopted mobilities two distinct factorizations serve all five compartments. The load vector carries all reaction terms at the previous time level, including $-vR$ in the R row, with the growth term switched to $-rS$ where $K < 0.01$, a branch the adjoint must inherit. A time step costs one vector assembly and one back-substitution.

6.5 The discrete adjoint of the IMEX scheme

We differentiate the discrete state system exactly (discretize-then-optimize) rather than discretizing the continuous adjoint, since only the former gives the exact gradient of the quantity we actually minimize. An inexact gradient can fail to be a descent direction near a stationary point, which stalls the line search.

Let $\mathbf{y}^n \in \mathbb{R}^{5N}$ collect all five compartments at $t^n = n\Delta t$ and let $\mathbf{v}^n \in \mathbb{R}^N$ be the control on $[t^n, t^{n+1})$, with $N_\mathrm{t} = T/\Delta t$ steps in total. The IMEX step then reads

$$\boldsymbol{\mathcal{A}}\,\mathbf{y}^{n+1} \;=\; \mathbf{L}\big(\mathbf{y}^n, \mathbf{v}^n\big), \qquad n = 0, \dots, N_\mathrm{t} - 1, \tag{46}$$

where $\mathbf{L}$ is the nonlinear load vector. Using a right-endpoint rule for the state terms and a left-endpoint rule for the control, the discrete objective is

$$J_h(\mathbf{v}) \;=\; \Delta t \sum_{n=1}^{N_\mathrm{t}} \int_{\Gamma_h} \big(c_1(S_{\infty,h} - S_h^n) + c_2 W_h^n\big)\,\mathrm{d}A \;+\; \Delta t \sum_{n=0}^{N_\mathrm{c}-1} \int_{\Gamma_h} \big[c_3 (v_h^n)^2 + c_4 v_h^n\big]\,\mathrm{d}A, \tag{47}$$

where $N_\mathrm{c} = T_\mathrm{c}/\Delta t$ is the number of controlled steps. The state sum runs over the whole horizon and the control sum over the control window, which is the continuation tail of Section 3.1 in discrete form. With one multiplier $\boldsymbol{\lambda}^{n+1} \in \mathbb{R}^{5N}$ per step, the discrete Lagrangian is

$$\mathcal{L}_h \;=\; J_h(\mathbf{v}) \;-\; \sum_{n=0}^{N_\mathrm{t}-1} \big(\boldsymbol{\lambda}^{n+1}\big)^{\mathsf{T}} \Big(\boldsymbol{\mathcal{A}}\,\mathbf{y}^{n+1} - \mathbf{L}(\mathbf{y}^n, \mathbf{v}^n)\Big). \tag{48}$$

Let $\mathbf{g} \in \mathbb{R}^{5N}$ be assembled from $\int_{\Gamma_h}(-c_1\psi_S + c_2\psi_W)\,\mathrm{d}A$, so that $\partial J_h/\partial \mathbf{y}^n = \Delta t\,\mathbf{g}$. Stationarity with respect to $\mathbf{y}^n$, $1 \le n \le N_\mathrm{t}$, then gives

$$\boldsymbol{\mathcal{A}}^{\mathsf{T}}\boldsymbol{\lambda}^{N_\mathrm{t}} = \Delta t\,\mathbf{g}, \qquad \boldsymbol{\mathcal{A}}^{\mathsf{T}}\boldsymbol{\lambda}^{n} \;=\; \left[\frac{\partial \mathbf{L}}{\partial \mathbf{y}}\big(\mathbf{y}^n, \mathbf{v}^n\big)\right]^{\mathsf{T}} \boldsymbol{\lambda}^{n+1} \;+\; \Delta t\,\mathbf{g}. \tag{49}$$

No term arises from $\mathbf{y}^0$, which does not depend on the control (Section 6.7). Since $\boldsymbol{\mathcal{A}}^{\mathsf{T}} = \boldsymbol{\mathcal{A}}$, the backward sweep reuses the forward factorizations. We assemble the transposed Jacobian by differentiating the load integrand with respect to each state, inheriting the $K < 0.01$ branch, and its terms are exactly the right-hand sides of (30)–(34) without their c_1, c_2 forcings.

Differentiating with respect to the control gives

$$\frac{\partial J_h}{\partial \mathbf{v}^n} = \Delta t\Big[\,2c_3\,\mathbf{M}_\mathrm{L}\,\mathbf{v}^n \;+\; c_4\,\mathbf{M}_\mathrm{L}\mathbf{1} \;-\; \mathbf{s}_v\big(\mathbf{y}^n, \boldsymbol{\lambda}^{n+1}\big)\Big], \qquad \big(\mathbf{s}_v\big)_j = \int_{\Gamma_h} R_h^n\,\lambda_{R,h}^{n+1}\,\varphi_j\,\mathrm{d}A, \tag{50}$$

for $0 \le n \le N_\mathrm{c} - 1$. The backward sweep still runs over the tail, which is how the tail's cost reaches $\boldsymbol{\lambda}^{N_\mathrm{c}}$, but no gradient is assembled there. Dividing by $\mathbf{M}_\mathrm{L}$ gives the Riesz representative $g_v^n \approx 2c_3 v^n + c_4 - R^n \lambda_R^{n+1}$, the discrete counterpart of (35), in the pairing

$$\big\langle a, b \big\rangle \;=\; \Delta t \sum_{n=0}^{N_\mathrm{c}-1} \sum_{i=1}^{N} \big(\mathbf{M}_\mathrm{L}\big)_{ii}\, a_i^n\, b_i^n, \qquad \big(\mathbf{M}_\mathrm{L}\big)_{ii} = \int_{\Gamma_h} \varphi_i\,\mathrm{d}A. \tag{51}$$

For block-constant controls, let $\beta(n)$ be the block containing step n, $N^{(b)}$ the number of steps in block b, and $\mathcal{E}$ the prolongation from block fields to step fields, zero past N_c. Representing the gradient of $\hat{J}_\mathrm{b}(z) := \hat{J}(\mathcal{E}z)$ in the pairing (51), restricted to block-constant fields, makes the block

gradient the *time average* $g_b = (N^{(b)})^{-1} \sum_{\beta(n)=b} g^n$, not the sum. This is the discrete counterpart of (38). The optimizer works in the block pairing with weights $\Delta t_b := N^{(b)} \Delta t$.

Finally, for (49) to be the exact transpose, every term of the adjoint load must use the same quadrature rule as the term of $\mathbf{L}$ it differentiates, or the Taylor test degrades from second order to first. Mass terms use the vertex rule that produces $\mathbf{M}_{\mathrm{L}}$, and all reaction, cost and gradient terms share one explicitly fixed Gauss degree. The quadrature of $\boldsymbol{\mathcal{A}}$ is unconstrained, since $\boldsymbol{\mathcal{A}}$ appears identically in both sweeps.

6.6 Spin-up: computing the initial condition

We compute the disease-free equilibrium S_∞ by running (43) from $S = K$, $E = I = R = W = 0$ until

$$\frac{\max_i \left| S_i^{n+1} - S_i^n \right|}{\Delta t \, \max_i K_i} \; < \; \varepsilon_{\mathrm{drift}} \tag{52}$$

or until a cap T_{spin} is reached, in which case the solver warns that the field has not converged. We take $\varepsilon_{\mathrm{drift}} = 10^{-5}$ yr^{-1} and $T_{\mathrm{spin}} = 60$ yr. The cap is generous because equilibration across a habitat patch of scale L is diffusive and takes a time of order L^2/κ.

6.7 Initial conditions

Each run starts from an established epizootic with all five compartments populated, rather than from an introduction. We build this initial state y^0 in two stages. First, we seed the disease on S_∞, since starting from $S = K$ would superimpose a demographic transient on the epidemic one at a comparable timescale ($1/r = 3.8$ yr). The seed is a Gaussian focus in horizontal coordinates,

$$\phi(\mathbf{x}) = \phi_0 \exp\Big(- \frac{\|\mathbf{x} - \mathbf{x}_0\|^2}{2\varsigma^2} \Big), \qquad \big(E, I, R\big)(\cdot, 0) = \big(\pi_E, \pi_I, \pi_R\big)\, \phi, \qquad W(\cdot, 0) = \frac{\rho_I \pi_I + \rho_R \pi_R}{\delta}\, \phi, \tag{53}$$

so all four disease compartments share one spatial profile. The amplitude ϕ_0 is the total infected density at the center $\mathbf{x}_0$, split by the stationary stage distribution $\sigma_1 E = \sigma_2 I = \alpha R$ into $\boldsymbol{\pi} = (0.338, 0.450, 0.212)$, with the reservoir at its $\dot{W} = 0$ level. We seed an established focus rather than a single case because 85% of R_0 runs through the environment, and a latent-only seed on a clean landscape would suppress most transmission while the reservoir charged. We use $\varsigma = 500$ m, the scale of a mule deer winter home range, and $\phi_0 = 2$ deer/km^2, which introduces about 3.1 infected animals. Both values, and the location of the focus, are [ASM].

Second, we advance the uncontrolled system from the seed, with the same step and $v \equiv 0$, until the accumulated clinical burden reaches a fixed detection threshold,

$$\int_0^{T_{\mathrm{intro}}} \int_\Gamma R \, \mathrm{d}\Gamma \, \mathrm{d}t \; \geq \; R_{\mathrm{det}}, \qquad R_{\mathrm{det}} = 60.8 \text{ clinical deer-years}, \tag{54}$$

and take all five fields at that moment as y^0, so the detection time T_{intro} is an output. Clinical animals are the ones most likely to be reported, sampled and found positive, so the rule approximates a fixed expected number of detections, although it ignores the pre-clinical infections that harvest testing also finds. Holding R_{det} fixed starts every run at the same accumulated clinical burden, so each run begins when its own surveillance would detect the disease. The value is a design constant rather than a field estimate. We set it to the burden accumulated in twenty years at $\delta = 0.25$ and a landscape-average R_0 of 2 on a 1397 km^2 test mesh of the Park Range in northern Colorado. Each run computes y^0 once and shares it between its baseline and its optimization.

6.8 Optimization

Each outer iteration consists of a forward sweep that stores the trajectory, a backward sweep that returns the gradient, and a step. The admissible set (18) is a box, and we minimize over it with the limited-memory, bound-constrained quasi-Newton method L-BFGS-B (Byrd et al., 1995; Nocedal and Wright, 2006), as implemented in SciPy. Since L-BFGS-B works in the Euclidean inner product of the vector it receives rather than in (51), we pass it $\boldsymbol{\zeta} = \mathbf{M}_{\mathrm{b}}^{1/2}\mathbf{v}$, where

$$\mathbf{M}_{\mathrm{b}} = \operatorname{diag}\big((\mathbf{M}_{\mathrm{L}})_{ii}\,\Delta t_b\big) \tag{55}$$

is the mass matrix of the block pairing. The Euclidean inner product of $\boldsymbol{\zeta}$ is then (51), the gradient with respect to $\boldsymbol{\zeta}$ is $\mathbf{M}_{\mathrm{b}}^{1/2}\nabla\hat{J}$, and the box stays a box. This scaling matters on the landscape meshes, since $(\mathbf{M}_{\mathrm{L}})_{ii}$ spans more than four orders of magnitude across a graded mesh.

6.9 Verification

Errors in an adjoint fail silently, since a wrong sign or a missing Jacobian term still gives plausible output. We verify the gradient with the Taylor test, in which the remainder $R_1(\varepsilon) = |\hat{J}(v + \varepsilon d) - \hat{J}(v) - \varepsilon\langle\nabla\hat{J}(v), d\rangle|$ decays as $O(\varepsilon^2)$ only if the gradient is exact, and we run it at a control strictly inside the bounds. Before this, a mesh-free check of the pointwise algebra compares the reaction terms, adjoint rows and gradient expression against numerically differentiated references at random states. It also verifies the structural identities: which compartments the control reaches, $S^* = K_0/R_0$ at the endemic equilibrium, and the exact time integral of (5). Over 600 random states its worst relative residual is 2.9×10^{-9}, at the truncation error of the central-difference reference. We run the Taylor test on the Upper Arkansas mesh of Section 7.5, where terrain and land cover exercise every term of the operator, along a full-support direction and along one supported on the last control block, which isolates the continuation tail. Both converge at rate 2 over five decades of ε (Table S5 of Online Resource 1).

7 Simulations

7.1 The sweep domain and the run design

We run the parameter sweep on a flat disk of radius 40 km with no landscape, meshed at a uniform $h = 250$ m (93,701 nodes). Results did not change qualitatively on meshes refined to $h = 100$ m. Here $c = p = 1$ everywhere, so $K = K_0 = 3$ deer/km^2, $\mathbf{D} = \kappa(\mathbf{e}_{\mathrm{s}} \otimes \mathbf{e}_{\mathrm{s}} + \mathbf{e}_{\mathrm{c}} \otimes \mathbf{e}_{\mathrm{c}})$ and $S_\infty = K_0$, giving a disease-free herd of 15,080 deer. The disease is seeded at the center, so every solution is radially symmetric, and with $c \equiv 1$ the best-habitat $\widetilde{R_0}$ of (13) is the local reproductive number, which we simply write as R_0 in this section. Removing the landscape lets us isolate how the swept parameters act, and Section 7.5 shows what $K(\mathbf{x})$ and $\mathbf{D}(\mathbf{x})$ add.

We vary $\delta \in \{0.125, 0.25, 0.5, 1\}$ yr^{-1} and $R_0 \in \{2, 3.5, 5\}$ as a full factorial of 12 cells, each with its own uncontrolled baseline, under the settings of Table 3. The R_0 axis runs from the low anchor of Section 2.9, where the infection spreads without depleting the herd behind it, to values where a depletion front crosses the disk. We also tested $R_0 = 10$, near the upper range of Almberg et al. (2011), but do not report it, since the herd collapses inside the window and control recovers only 1–4% of $\hat{J}$. The control horizon $T_{\mathrm{c}} = 20$ yr is about the shortest window over which a culling policy can be evaluated, since CWD epizootics take one to three decades to reach double-digit prevalence, and its annual blocks match the period on which agencies set quotas. Each cell starts from its own detection state (Section 6.7), reached between 5.5 yr at ($\delta = 1,\ R_0 = 5$) and 19.1 yr

at ($\delta = 0.125,\ R_0 = 2$) after seeding, at a domain prevalence of 0.24–1.98%. Every optimization stopped on the relative-reduction criterion in 8 to 10 iterations, with the projected-gradient residual (40) reduced by a factor of at least 8.6×10^3. This certifies stationarity, not global optimality.

The grid sets $c_2 = 0$ because pricing the reservoir would confound the δ sweep. Since $W^* \propto 1/\delta$ while the host equilibrium does not depend on δ, a fixed c_1:c_2 would weight the reservoir less and less as δ rises. Section 7.3 checks that this choice does not drive the placement of effort.

7.2 The uncontrolled wave

Figure 2 shows the radial structure of the uncontrolled invasion at the slowest and fastest reservoirs. Transmissibility sets how deeply the herd is depleted behind the front, and reservoir persistence sets how soon the infection leaves the depleted core. At $R_0 = 2$ the infection is a dome centered on the focus, where the herd still holds 61–82% of S_∞, and the three host classes share nearly one profile. At higher R_0 the core is depleted to 3–33% of S_∞, and the infection leaves it as an outward pulse. A fast reservoir empties the core quickly, so the pulse forms within six years of handover, while a slow one keeps the core infectious with prions shed early in the epizootic. At ($\delta = 0.125,\ R_0 = 3.5$), for example, the infected density is still highest at the focus in year 12. Within a pulse the classes are ordered by stage, with the exposed peak 2.5–5.0 km ahead of the clinical one and the reservoir behind both.

The front also speeds up with both parameters. Its leading edge, the outermost radius at which $E + I + R \geq 0.1$ deer/km^2, advances at 0.58 km/yr at ($\delta = 0.125,\ R_0 = 2$) and at 2.94 km/yr at ($\delta = 1,\ R_0 = 5$), and reaches the rim within the control horizon at eight of twelve cells. Whether the herd collapses depends on R_0. At $R_0 = 2$ the herd declines by 6 to 48% over the full 30 yr and the epizootic is still growing at T, so every reduction we report there is in an integral over a growing trajectory. At $R_0 \geq 3.5$ the epizootic burns through the herd within the window, reducing it by 76–94% everywhere except ($\delta = 0.125,\ R_0 = 3.5$). Table S6 of Online Resource 1 gives each cell's handover state and uncontrolled outcome.

7.3 Where the effort goes

Figure 3 sets the optimal control of one cell against its state. We measure positions from the leading edge, against the state averaged over each control block, and average them over the years in which the edge lies inside 38 km. Where the herd is not depleted, the effort forms a disk around the focus. At $R_0 = 2$ and $\delta \leq 0.25$ it peaks at the focus, where the clinical density and RS also peak, and falls to zero 10–22 km beyond the leading edge. At $\delta \geq 0.5$ it flattens into a plateau, marking the transition to a depletion front.

Where a depletion front forms, the effort meets the invasion at its leading edge. At $R_0 \geq 3.5$ the control is a ring that travels outward with the invasion, with its peak on average 0.4–3.8 km inside the leading edge. This places the peak 10–13 km outside the maximum of RS, except at ($\delta = 1,\ R_0 = 5$) where the gap is 6.2 km, and 12–19 km outside the clinical peak. Since the optimality condition places effort where $\lambda_R R$ exceeds $c_4 + 2c_3 v$ (Section 5), a peak this far out means that λ_R rises steeply across the front, toward the intact herd that the next clinical animal would reach, so the effort follows the adjoint rather than the clinical field. The placement does not come from the S-gating of the deficit source (Section 4.2). Moving the whole state weight onto W, with c_2 scaled so that the uncontrolled J is unchanged, leaves the peak at the edge and changes total effort by 2% (Fig. 3k–o), since the source of W reaches λ_R without passing through λ_S. Toward the end of the window the ring falls back to 5–7.5 km behind the leading edge, which may reflect the approaching rim or the end of the control at T_c.

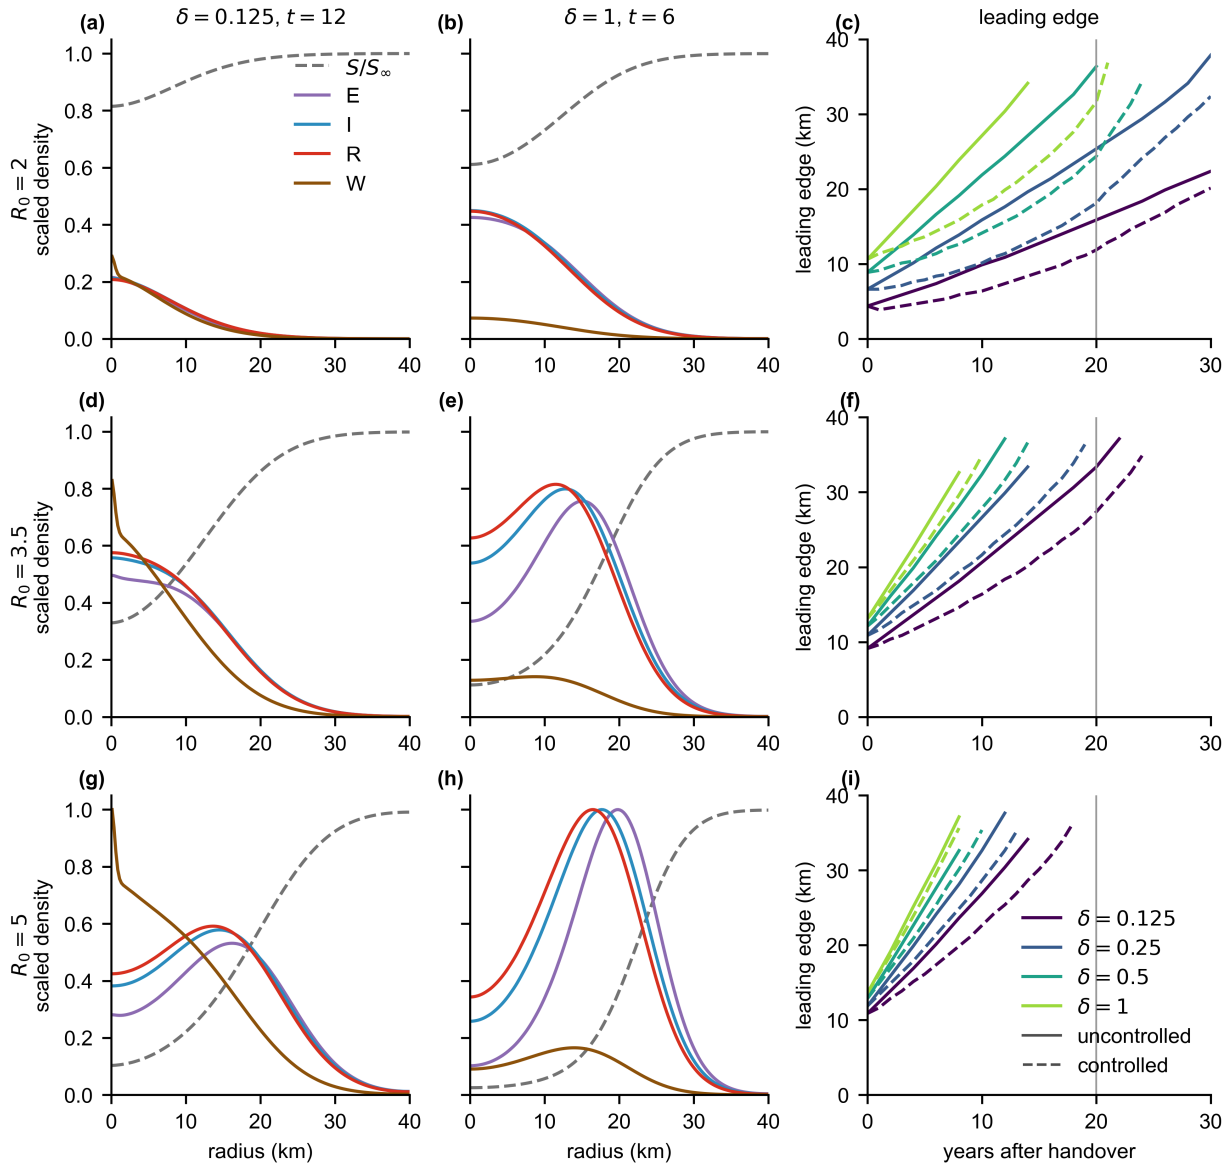


Fig. 2 The invasion on the flat disk, with $R_0 = 2$ in (a)–(c), 3.5 in (d)–(f) and 5 in (g)–(i). Uncontrolled radial profiles twelve years after handover at $\delta = 0.125$ (a, d, g) and six years after at $\delta = 1$ (b, e, h). S/S_∞ is dashed, and E, I, R and W are each scaled by their largest value over these six panels, so the panels compare magnitudes as well as shapes. Panels (c), (f) and (i) show the leading edge, the outermost radius at which $E + I + R \geq 0.1$ deer/km^2, against time for every δ, uncontrolled (solid) and under the optimal control (dashed). Each curve stops when the edge reaches the rim or the infection falls below the threshold everywhere. The gray line marks T_c

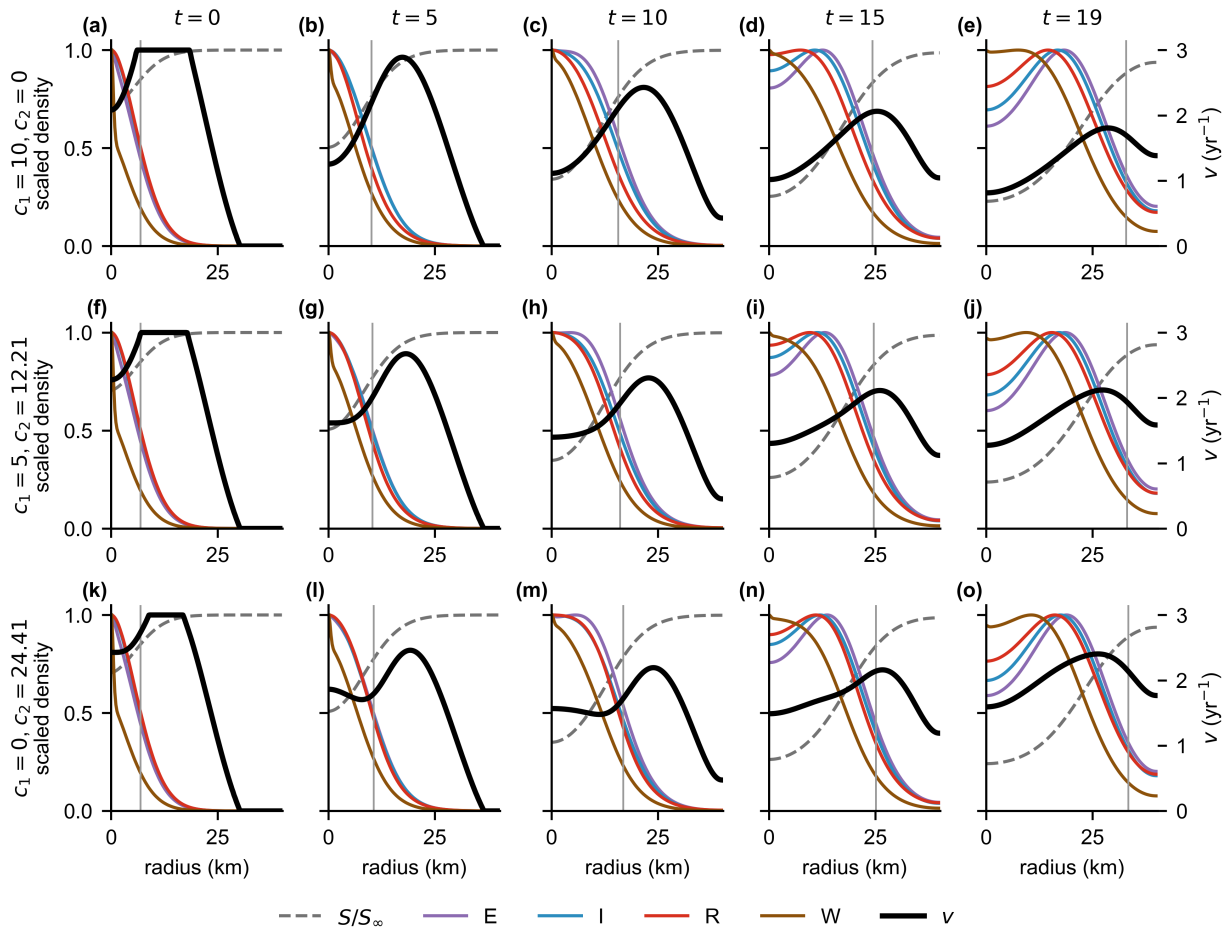


Fig. 3 Radial state and control at ($\delta = 0.25$, $R_0 = 3.5$) under three objectives. Panels (a)–(e) are the reference $c_1 = 10$, $c_2 = 0$. Below them the state weight moves onto W, to $(c_1, c_2) = (5, 12.21)$ in (f)–(j) and $(0, 24.41)$ in (k)–(o), with c_2 scaled so that the uncontrolled J is unchanged. The panel at $t = s$ shows the state at s and the control of the block $[s, s+1)$. The front advances about 1.6 km over a block, so against the block-mean state used in the text the effort sits about 0.8 km further in than drawn. Each density is scaled by its own maximum in the panel. S/S_∞ is dashed and v (thick, right axis) is on one scale shared by all panels. The gray line marks the radius at which the infected profile falls to half its plateau. In every row the effort peak runs ahead of the clinical peak. Pricing W lowers the ring and raises the core without moving the peak

The depleted core is not abandoned. Behind the front the effort falls but never reaches zero, and instead it tracks the susceptibles. At ($\delta = 1$, $R_0 = 3.5$) it falls from 1.59 yr^{-1} at handover to 0.13 in year 6, when the core holds 18% of its disease-free herd, and climbs back to 1.24 by year 19 as the herd recovers. The zero set created by c_4 lies ahead of the ring rather than in the core.

7.4 What the control achieves

The optimal control slows the invasion but does not stop it (Fig. 2c, f, i and Table 4). In every cell the controlled front keeps advancing through the horizon, and once the control ends at T_c it accelerates again. At $R_0 = 2$ the control holds the front to 0.56–0.68 of its uncontrolled speed and

delays its arrival at 20 km by 3.6–7.9 yr. At $R_0 \geq 3.5$ the front keeps 75–94% of its uncontrolled speed, and the delay shrinks as δ rises, from 4.4 to 0.7 yr at $R_0 = 3.5$ and from 2.2 to 0.3 yr at $R_0 = 5$.

Table 4 The optimal control on the flat disk, as percentage change against each cell's own uncontrolled baseline. "Hazard" is the mean force of infection (56), "mean v" the mass- and block-weighted realized effort over the disk, "speed" the controlled leading-edge speed over the uncontrolled one while the edge is interior, and "delay" how much later the controlled edge reaches 20 km. Table S7 of Online Resource 1 gives further measures.

R_0	δ	Deficit	Hazard	Mean v	Speed	Delay (yr)
2	0.125	−28.6	−27.1	0.19	0.68	3.6
	0.25	−45.3	−40.8	0.55	0.61	6.7
	0.5	−53.3	−46.2	1.15	0.56	7.9
	1.0	−51.1	−40.1	1.66	0.60	6.4
3.5	0.125	−36.8	−30.9	0.96	0.75	4.4
	0.25	−30.7	−21.3	1.46	0.81	2.7
	0.5	−18.7	−9.6	1.41	0.83	1.4
	1.0	−10.0	−0.5	1.02	0.88	0.7
5	0.125	−25.7	−17.6	1.37	0.83	2.2
	0.25	−15.4	−9.5	1.29	0.83	1.2
	0.5	−9.2	−3.3	1.05	0.90	0.6
	1.0	−7.4	+2.6	0.83	0.94	0.3

The accumulated deficit, the objective's own measure, follows the same pattern. The control recovers 29–53% of it at $R_0 = 2$ against 7–26% at $R_0 = 5$. At $R_0 \geq 3.5$ the benefit falls monotonically with δ, and where the herd collapses fastest the control recovers a tenth of the deficit or less. At $R_0 = 2$ the benefit instead peaks at $\delta = 0.5$, since a weak epizootic with a persistent reservoir barely grows inside the window and leaves little to prevent. Part of why a high R_0 is harder to control is that its epizootic is further along when it is detected. At every δ, detection at $R_0 = 5$ hands over about three times as many infected animals as at $R_0 = 2$.

We measure the per-deer risk of infection by the S-weighted mean force of infection,

$$\overline{\Lambda} \;=\; \frac{\displaystyle\int_0^T \!\!\int_\Gamma S\,\Lambda \,\mathrm{d}\Gamma\,\mathrm{d}t}{\displaystyle\int_0^T \!\!\int_\Gamma S \,\mathrm{d}\Gamma\,\mathrm{d}t}, \tag{56}$$

quoted per 1000 susceptible deer-years, rather than by cumulative incidence. Incidence rises under control at 5 of 12 cells, by up to 17%, because the control keeps deer alive and so exposes more of them. On the per-deer measure control helps at 11 of 12 cells, and changes it by only −0.5 and +2.6% at the two $\delta = 1$ cells with $R_0 \geq 3.5$, where the control buys standing deer and little else.

7.5 What the landscape adds

We now restore the landscape on three 100×100 km Colorado meshes, to show what $K(\mathbf{x})$ and $\mathbf{D}(\mathbf{x})$ add rather than to forecast these regions. Fitting the model to a region would mean replacing κ, K_0 and the land-cover coefficients with local estimates (Section 8). In the Upper Arkansas (Fig. 4) and around Carbondale, terrain carves the habitat into branching valleys, whereas around Estes

Park (Fig. 5) the main contrast is in the land cover, between the foothills and the plains. Terrain is the USGS 3DEP 1/3 arc-second DEM (U.S. Geological Survey, 2024a) and land cover the 2025 Annual NLCD (U.S. Geological Survey, 2024b). Section S3 of Online Resource 1 describes the three meshes and gives the Carbondale figures. Each mesh is run at $\delta = 0.25$ with a landscape-average R_0 of 3.5, so the disk's ($\delta = 0.25$, $R_0 = 3.5$) cell serves as the reference. On each mesh we cross the slope tensor with the land cover in a 2×2 ablation. The isotropic arms set $\eta \equiv 1$ at unchanged κ, and the flattened arms replace p and c by their area-weighted means. Since κ is not rescaled by the mean permeability $\overline{p} = 0.70$–0.78, every landscape invasion runs slower than the disk's, so the comparison with the disk is qualitative.

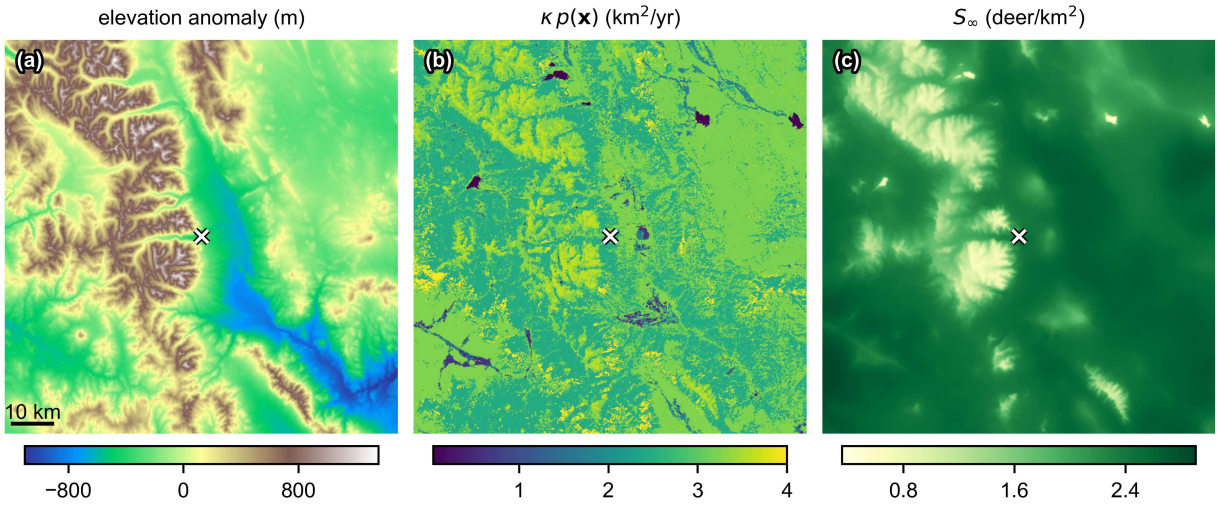


Fig. 4 The Upper Arkansas mesh: (a) terrain as the anomaly about its mean, (b) the land-cover diffusivity $\kappa\, p(\mathbf{x})$ and (c) the disease-free equilibrium S_∞. The cross marks the focus. S_∞ is highest on the valley floors and lowest on the ridges

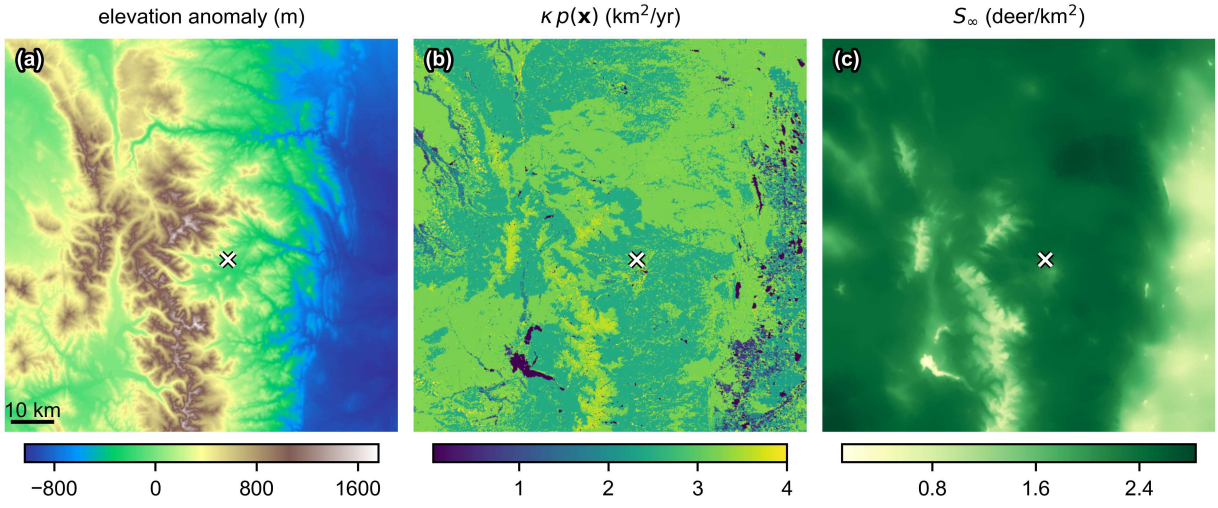


Fig. 5 The Estes Park mesh, on the same terms as Fig. 4. S_∞ is lowest on the high peaks to the west and on the developed land and cropland of the plains to the east

The slope tensor already shapes S_∞, which is K smoothed by the diffusion operator. The smoothing length is $\sqrt{\kappa\eta/r}$ up the slope and $\sqrt{\kappa(2-\eta)/r}$ along the contour, so it goes from 3.92 km in both directions on flat ground to 0.39 and 5.5 km at 45°. The tensor therefore preserves habitat boundaries on steep ground that isotropic diffusion erases, and above 600 m/km of relief it sharpens $|\nabla_\Gamma S_\infty|$ about fourfold.

The landscape barely moves the benefit. On each mesh the deficit reduction varies by at most 3.7 points across the four arms, and every run recovers 28–34% of the deficit, against 30.7% at the matching cell on the disk (Table S4 of Online Resource 1). What the landscape changes is where the effort goes. The full model reaches its benefit with the least effort on every mesh, 26–35% less than with neither mechanism, and each mechanism contributes on its own. Most of this saving is ground the invasion barely reaches, which the optimal policy leaves untreated (Table S4 of Online Resource 1).

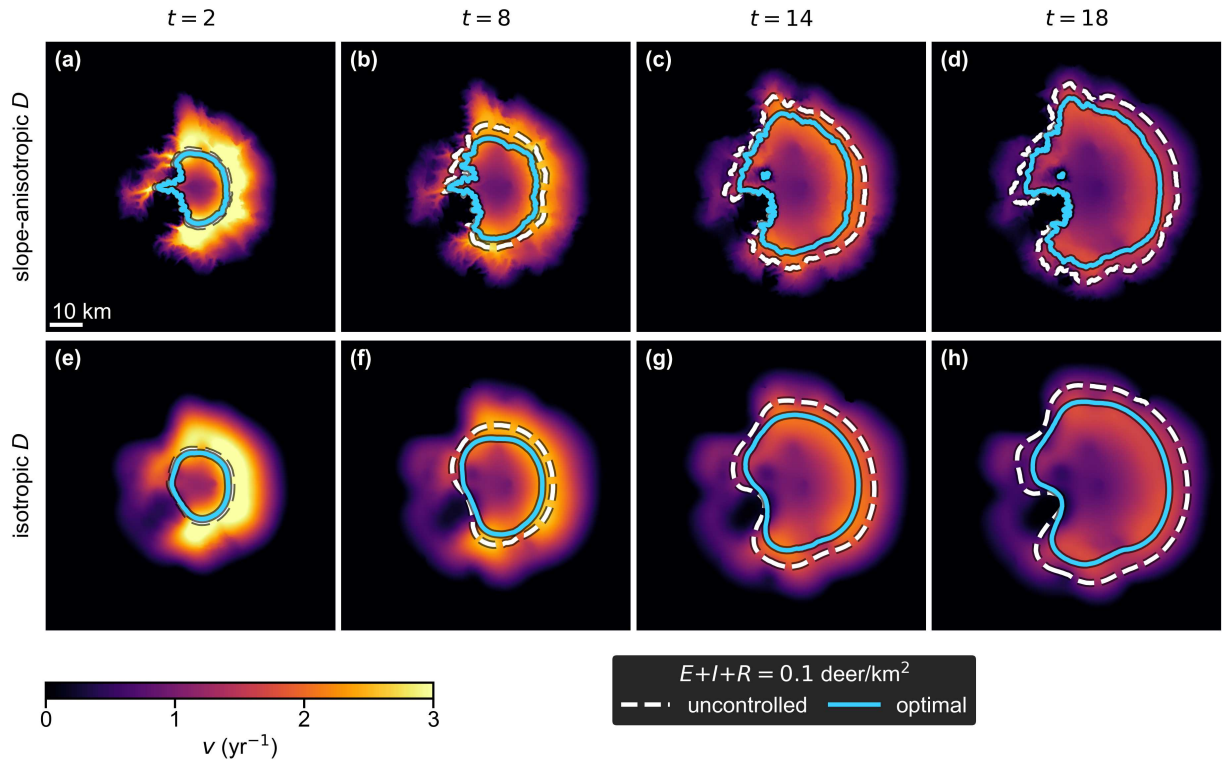


Fig. 6 The Upper Arkansas at $\delta = 0.25$, $\widetilde{R_0} = 4.24$, $c_4 = 0.5$, with the land cover measured, at $t = 2$, 8, 14 and 18. Panels (a)–(d) show the optimal v (inferno, $[0, v_{\max}]$) under the slope-anisotropic $\mathbf{D}$ and (e)–(h) under the isotropic one. The panel at $t = s$ shows the fronts at s and the control of the block $[s, s+1)$. The dashed white line is the uncontrolled 0.1 deer/km^2 front and the solid cyan line the controlled front. By $t = 2$ the front already runs west in narrow fingers up the tributary valleys. Table S4 of Online Resource 1 summarizes all four arms

Figures 6 and 7 set each measured-land-cover arm's optimal control against its uncontrolled front. As on the disk, the control forms a ring at the invasion front and holds the front several kilometers behind the uncontrolled one by year 20. On the Upper Arkansas the tensor breaks the western front into fingers along the tributary valleys and the control follows them, whereas isotropic diffusion merges them into one lobe. The heaviest effort, however, is placed by the land

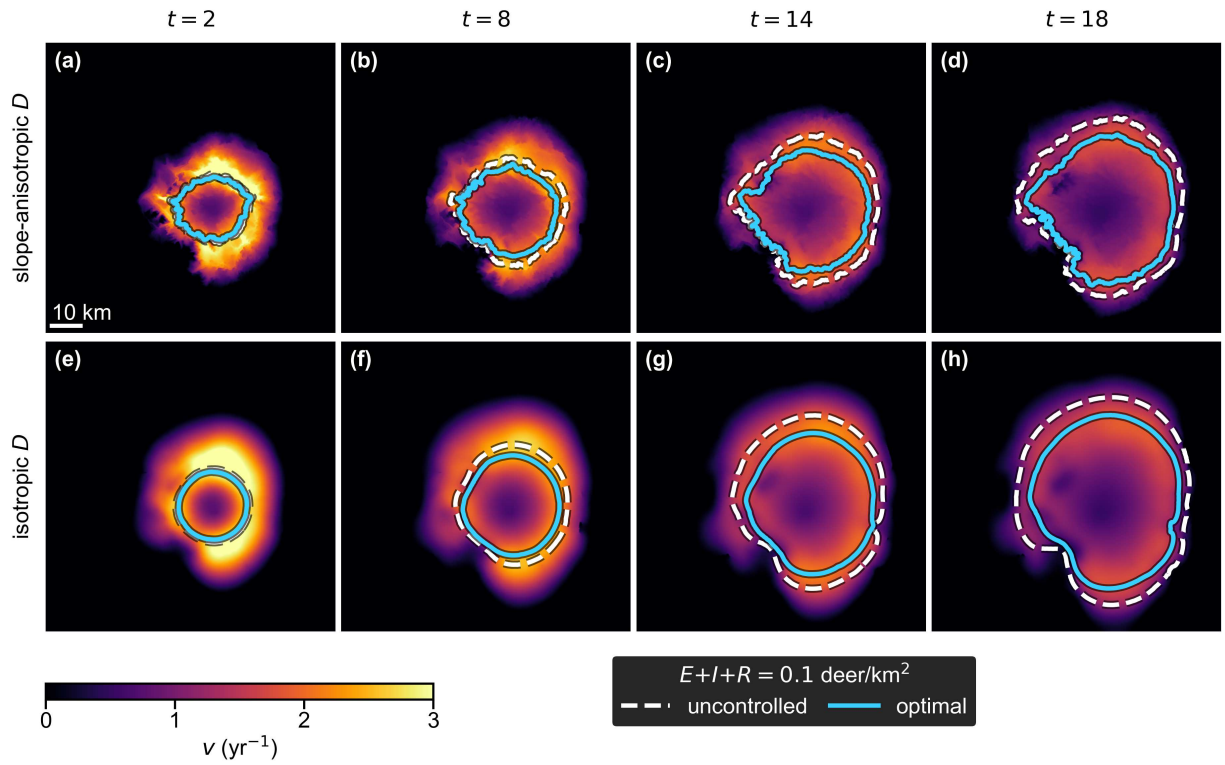


Fig. 7 The Estes Park mesh at $\widetilde{R_0} = 4.50$, on the same terms as Fig. 6. The front is nearly round at $t = 2$

cover rather than by the slope. In the terrain-driven valleys the land cover sends effort to the low ground, and the tensor strengthens this. On the Upper Arkansas, effort is 3.2 times heavier in the lowest elevation quartile than in the highest in the full model, against 0.9 with neither mechanism.

7.6 Robustness

The control rarely saturates. The space–time share at $v_{\max}$ is at most 4.3%, so no reduction in Table 4 is limited by the cap. The magnitudes do rest on the two transmission ratios we know least well (Section 3.2). Rerunning the ($\delta = 0.25$, $R_0 = 3.5$) cell at $\varrho = 5$, and at $m = 2.6$ and 18.2, moves the recovered deficit between 20 and 37% about the reference 31%. In every arm, however, $R_c(v_{\max}) \geq 2.31$ and the front keeps 74–82% of its uncontrolled speed, so the qualitative conclusions hold. On the Park Range test mesh (Section 6.7), doubling the continuation tail moved the reported reduction by 1.3 points, and refining Δt to 1/40 yr changed the uncontrolled objective by 0.17%. Section S4 of Online Resource 1 gives these checks in full.

8 Discussion

8.1 What a sight-based cull can and cannot do

We deliberately kept the formulation narrow. Splitting the shedding population into pre-clinical and clinical stages lets us aim the control at animals a field crew can see, and shows that this visible class carries about 60% of transmission, mostly through the carcass and the late-stage shedding

that the reservoir accumulates. The optimizer has no direct lever on the remaining 40%, and every infection it acts on has already been transmitting, unseen, for about forty weeks. Despite this blind spot, the control slows the front by up to 44% and recovers 7–53% of the accumulated susceptible deficit, with transmissibility deciding most of the difference. It never halts the front or eliminates the disease, as (21) predicts before any computation.

Since this bound does not depend on the spatial model, field estimates of ϱ and m, the two ratios we know least well, would show whether such a cull is worth attempting before any spatial model is fitted (Section 3.2). The results are stationary points of a nonconvex problem under stated weights rather than certified optima, and where the epizootic does not turn over inside the window, they are reductions in an integral over a growing trajectory.

The epizootic crosses the domain as a front, and local dynamics relax toward their endemic level on e-folding times of 14 yr or longer (Section 2.10), so nothing in the runs is near a fixed point. The model describes the first few decades of an epizootic on a real landscape, which is the window over which a policy is chosen and that surveillance covers. Its well-mixed endemic limits are properties of the kinetics rather than predictions.

8.2 Implications for management

Three findings bear directly on how culling effort might be allocated, with the caveat that the model has not been fitted to any specific region. First, where transmission is strong enough to deplete the herd, the optimal effort is concentrated at the leading edge of the invasion rather than at the clinical peak or the point of first detection (Section 7.3). Effort is worth most where the next clinical animals would reach an intact herd. This is consistent with the field result of Mandujano Reyes et al. (2025) that targeted culling in Wisconsin reduced the spatial spread of CWD rather than its local growth rate. Miller Neilan and Lenhart (2011) found a similar placement for rabies in raccoons, where the optimal vaccine distribution is an arc of bait ahead of the infectious wave, although vaccination protects susceptibles directly while a sight-based cull removes only infected animals. Second, the benefit of a fixed effort falls steeply with transmissibility and, at moderate to high transmissibility, with the speed of the reservoir (Table 4), so early detection and an estimate of R_0 largely set what a culling program can expect to achieve. Third, cumulative incidence is a poor metric for evaluating such a program, since a cull that keeps more deer alive also exposes more of them. Per-capita measures such as (56) or the herd deficit avoid this artifact (Section 7.4).

On real terrain, the landscape changes the policy far more than its benefit. Terrain and land cover together let the optimizer reach nearly the same benefit with a quarter to a third less effort, mainly because poor habitat and steep ground carry less of the invasion and need no culling (Section 7.5). Miller Neilan and Lenhart (2011) found the same effect for rabies, where a river let a partial arc of bait do the work of the full arc needed on a uniform domain. This saving compares optimal policies across models, however, and not policies on one landscape.

8.3 Limitations

The model makes several structural simplifications. There is no background mortality, so α absorbs disease-induced death together with any predation or harvest that selectively takes symptomatic animals, v is total removal effort rather than effort above a normal hunt, and $v \equiv 0$ is an unmanaged herd. As a result, the optimizer counts every averted infection as a certain future shedder, which biases optimal effort upward. At the annual adult survival of 0.82 reported by Geremia et al. (2015) and Miller et al. (2008), about a tenth of infected animals would die before shedding and over a fifth before becoming clinical. Adding a background mortality of $\mu = 0.198$ yr^{-1} would require re-

solving the transmission coefficients against R_0. The model also has no age or sex structure, which caps endemic prevalence at 31% (Section S1 of Online Resource 1), and no seasonal migration. We read Γ as winter range and every parameter as an annual average over the months the herd occupies it, so nothing here addresses summer-range prevalence or the timing of a cull within a year. Directed relocation would need an advection term, which is not self-adjoint and would cost us the shared factorization, so a seasonal signal in β_1 and β_2 is a cheaper alternative.

The effort cap $v_{\max}$ limits a per-animal removal rate, whereas in practice the binding constraint is search. On the disk, clinical density at handover is one animal per 14–42 km^2 of invaded ground, so effort scales with area covered rather than animals found, and a cap on effort per unit area would be the more faithful constraint. Since a generous $v_{\max}$ flatters the positive results and strengthens the negative ones, we hold it fixed and report saturation. A continuum model also has no extinction threshold, so a residual density can regrow a front during the tail, and we make no claim of local elimination.

Two inputs carry most of the remaining uncertainty. The first is the pair of transmission ratios ϱ and m, which set the 60:40 split and so the size of the blind spot. Across the arms of Section 7.6 the recovered deficit ranges from 20 to 37% about the reference 31%, while $R_c(v_{\max})$ stays above 2.3, so the qualitative conclusions hold across the bracket but the magnitudes do not. Of the two, ϱ is the better target for field estimation, since it only shifts shedding later within an infection and so changes mainly what the cull can see, while m also routes transmission through W, where it waits about $1/\delta$, and so changes the speed of the epizootic as well. The second, the diffusivity κ, matters less than it first appears. On the disk it only sets the length scale, so the results there hold at any κ once distances are rescaled (Section 2.5). On the landscapes it sets the diffusion length $\sqrt{\kappa/r}$ against the width of valleys and habitat patches, and a larger value would smooth more of the land cover out of S_∞, so the results of Section 7.5 are conditional on it.

8.4 Extensions

On the landscapes, κ cannot be raised at fixed domain size without driving the invasion into the boundary, and the natural follow-up is to enlarge κ and the domain together using the ecological diffusion model of Garlick et al. (2014) (Section 2.5). That model uses GPS-fitted motilities for each cover class in place of a fitted scalar, and aggregates the population in good habitat without a separate carrying-capacity field, so it would replace both of our weakest inputs at once. The transmission coefficients should also be reconciled with the fitted estimates of Wasserberg et al. (2009) and Miller et al. (2006), and K_0 and the land-cover coefficients calibrated against regional surveillance and density data.

The most interesting extensions of the control address what v cannot reach. A second control on the reservoir, through soil treatment, would enter (5) as $-v_2 W$. This would be a qualitative addition, since the current model can prevent contamination but not remove it. Removing the carcasses of animals that die naturally is a separate lever that the model does not include. It acts on deposition rather than on W, is not limited to the clinical window, and averts a pulse that carries about 41% of R_0. That share rests on reading 49% of each animal's deposit, the 100 of 206 prion units of Almberg et al. (2011), as a removable carcass (Section 2.7), and shrinks if less of the deposit is tied to the carcass. A non-selective cull that reaches the pre-clinical class is the other half of the gap, and the deficit running cost already prices the healthy animals such a cull would remove. Smaller extensions include a take limit $\iint vR \le C$ on animals removed, restriction to accessible management units, and imperfect detection.

9 Conclusions

We have developed a spatially explicit model of chronic wasting disease on real terrain and used it as the constraint of an optimal control problem for a cull restricted to clinically affected deer. A closed-form bound shows that, under a plausible cap on effort, such a cull cannot by itself bring transmission below replacement anywhere in our range $2 \leq R_0 \leq 5$, because it never reaches the pre-clinical animals that generate about 40% of transmission. Within that limit, the exact discrete adjoint lets us compute where and when a bounded effort does the most good. Optimal effort meets the invasion at its leading edge, slows the front to 56–94% of its uncontrolled speed and recovers 7–53% of the accumulated loss of susceptible deer, with transmissibility deciding most of the difference. On real landscapes, terrain and land cover change where effort is needed far more than how much can be achieved. The analysis identifies the deposition ratio ϱ and the route ratio m as the quantities whose field estimation would most sharpen these conclusions, and points to carcass removal and environmental decontamination as the levers most likely to close the gap a sight-based cull leaves open.

Statements and Declarations

Funding

No funding was received for conducting this study.

Competing interests

The authors have no competing interests to declare that are relevant to the content of this article.

Ethics approval

Not applicable. This study uses only published parameter values and public elevation and land-cover data.

Author contributions

Gabe Flanagin: conceptualization, model formulation, parameter identification, adjoint derivation, finite element formulation and implementation, mesh generation, discrete adjoint and optimization, data curation, simulation design and analysis, visualization, and writing of the manuscript. Alex Escoto: model formulation, adjoint derivation, finite element formulation and implementation, and mesh generation. Zhen Chao: advising and review of the manuscript.

Code and data availability

The source code is available at `https://github.com/gflanagin/CWD_codebase`, and the code, mesh bundles, grid output and figure scripts are archived at `https://doi.org/10.5281/zenodo.23093270`. The implementation consists of five files. `cwd_control_problem.py` builds the mixed space and land-cover fields, assembles and factorizes $\boldsymbol{\mathcal{A}}$, and performs the forward and backward sweeps. `CWD_optimal_control.py` is the driver and implements the Taylor test under `--gradient-check`. `utilities/check_reaction_adjoint.py` is the mesh-free check of Section 6.9, and the two spin-up utilities build the initial condition of Section 6.7. A table mapping every symbol of Sections 3–6.5 to the corresponding variational form or configuration key accompanies the source. Elevation and

land-cover data are publicly available from the U.S. Geological Survey (U.S. Geological Survey, 2024a,b), and the mule deer movement data used in Section 2.6 are available from Barrile et al. (2024b).

References

Emily S. Almberg, Paul C. Cross, Christopher J. Johnson, Dennis M. Heisey, and Bryan J. Richards. Modeling routes of chronic wasting disease transmission: environmental prion persistence promotes deer population decline and extinction. *PLoS ONE*, 6(5):e19896, 2011. https://doi.org/10.1371/journal.pone.0019896.

Gabriel M. Barrile, Paul C. Cross, Cody Stewart, Jennifer Malmberg, Rhiannon P. Jakopak, Justin Binfet, Kevin L. Monteith, Brandon Werner, Jessica Jennings-Gaines, and Jerod A. Merkle. Chronic wasting disease alters the movement behavior and habitat use of mule deer during clinical stages of infection. *Ecology and Evolution*, 14(5):e11418, 2024a. https://doi.org/10.1002/ece3.11418.

Gabriel M. Barrile, Paul C. Cross, Cody Stewart, Jennifer Malmberg, Rhiannon P. Jakopak, Justin Binfet, Kevin L. Monteith, Brandon Werner, Jessica Jennings-Gaines, and Jerod A. Merkle. Chronic wasting disease alters the movement behavior and habitat use of mule deer during clinical stages of infection. Dryad dataset, 2024b.

Richard H. Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, 1995. https://doi.org/10.1137/0916069.

Timothy Clayton, Scott Duke-Sylvester, Louis J. Gross, Suzanne Lenhart, and Leslie A. Real. Optimal control of a rabies epidemic model with a birth pulse. *Journal of Biological Dynamics*, 4(1):43–58, 2010. https://doi.org/10.1080/17513750902935208.

Melia T. DeVivo, David R. Edmunds, Matthew J. Kauffman, Brant A. Schumaker, Justin Binfet, Terry J. Kreeger, Bryan J. Richards, Hermann M. Schätzl, and Todd E. Cornish. Endemic chronic wasting disease causes mule deer population decline in Wyoming. *PLoS ONE*, 12(10): e0186512, 2017. https://doi.org/10.1371/journal.pone.0186512.

David R. Edmunds, Matthew J. Kauffman, Brant A. Schumaker, Frederick G. Lindzey, Walter E. Cook, Terry J. Kreeger, Ronald G. Grogan, and Todd E. Cornish. Chronic wasting disease drives population decline of white-tailed deer. *PLoS ONE*, 11(8):e0161127, 2016. https://doi.org/10.1371/journal.pone.0161127.

Luis E. Escobar, Sandra Pritzkow, Steven N. Winter, Daniel A. Grear, Megan S. Kirchgessner, Ernesto Dominguez-Villegas, Gustavo Machado, A. Townsend Peterson, and Claudio Soto. The ecology of chronic wasting disease in wildlife. *Biological Reviews*, 95(2):393–408, 2020. https://doi.org/10.1111/brv.12568.

Cody E. FitzGerald and James P. Keener. The effect of predation on the dynamics of chronic wasting disease in deer. *arXiv preprint arXiv:2509.00134*, 2025. https://doi.org/10.48550/arXiv.2509.00134.

Martha J. Garlick, James A. Powell, Mevin B. Hooten, and Leslie R. MacFarlane. Homogenization, sex, and differential motility predict spread of chronic wasting disease in mule deer in

southern Utah. *Journal of Mathematical Biology*, 69(2):369–399, 2014. https://doi.org/10.1007/s00285-013-0709-z.

Chris Geremia, Michael W. Miller, Jennifer A. Hoeting, Michael F. Antolin, and N. Thompson Hobbs. Bayesian modeling of prion disease dynamics in mule deer using population monitoring and capture-recapture data. *PLoS ONE*, 10(10):e0140687, 2015. https://doi.org/10.1371/journal.pone.0140687.

Trevor J. Hefley, Mevin B. Hooten, Robin E. Russell, Daniel P. Walsh, and James A. Powell. When mechanism matters: Bayesian forecasting using models of ecological diffusion. *Ecology Letters*, 20(5):640–650, 2017. https://doi.org/10.1111/ele.12763.

Michael Hinze, René Pinnau, Michael Ulbrich, and Stefan Ulbrich. *Optimization with PDE Constraints*, volume 23 of *Mathematical Modelling: Theory and Applications*. Springer, Dordrecht, 2009. https://doi.org/10.1007/978-1-4020-8839-1.

Christopher J. Johnson, Kristen E. Phillips, Peter T. Schramm, Debbie McKenzie, Judd M. Aiken, and Joel A. Pedersen. Prions adhere to soil minerals and remain infectious. *PLoS Pathogens*, 2 (4):e32, 2006. https://doi.org/10.1371/journal.ppat.0020032.

Suzanne Lenhart and John T. Workman. *Optimal Control Applied to Biological Models*. Chapman and Hall/CRC, Boca Raton, FL, 2007. https://doi.org/10.1201/9781420011418.

Juan Francisco Mandujano Reyes, Ting Fung Ma, Ian P. McGahan, Daniel J. Storm, Daniel P. Walsh, and Jun Zhu. Spatiotemporal causal inference with mechanistic ecological models: evaluating targeted culling on chronic wasting disease dynamics in cervids. *Environmetrics*, 36(2): e2901, 2025. https://doi.org/10.1002/env.2901.

Candace K. Mathiason, Jenny G. Powers, Sallie J. Dahmes, David A. Osborn, Karl V. Miller, Robert J. Warren, Gary L. Mason, Sheila A. Hays, Jeanette Hayes-Klug, Davis M. Seelig, Margaret A. Wild, Lisa L. Wolfe, Terry R. Spraker, Michael W. Miller, Christina J. Sigurdson, Glenn C. Telling, and Edward A. Hoover. Infectious prions in the saliva and blood of deer with chronic wasting disease. *Science*, 314(5796):133–136, 2006. https://doi.org/10.1126/science.1132661.

Michael W. Miller, Elizabeth S. Williams, N. Thompson Hobbs, and Lisa L. Wolfe. Environmental sources of prion transmission in mule deer. *Emerging Infectious Diseases*, 10(6):1003–1006, 2004. https://doi.org/10.3201/eid1006.040010.

Michael W. Miller, N. Thompson Hobbs, and Simon J. Tavener. Dynamics of prion disease transmission in mule deer. *Ecological Applications*, 16(6):2208–2214, 2006. https://doi.org/10.1890/1051-0761(2006)016[2208:DOPDTI]2.0.CO;2.

Michael W. Miller, Heather M. Swanson, Lisa L. Wolfe, Fred G. Quartarone, Sherri L. Huwer, Charles H. Southwick, and Paul M. Lukacs. Lions and prions and deer demise. *PLoS ONE*, 3 (12):e4019, 2008. https://doi.org/10.1371/journal.pone.0004019.

Rachael Miller Neilan and Suzanne Lenhart. Optimal vaccine distribution in a spatiotemporal epidemic model with an application to rabies and raccoons. *Journal of Mathematical Analysis and Applications*, 378(2):603–619, 2011. https://doi.org/10.1016/j.jmaa.2010.12.035.

Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, New York, 2 edition, 2006. https://doi.org/10.1007/978-0-387-40065-5.

Katherine L. Parker, Charles T. Robbins, and Thomas A. Hanley. Energy expenditures for locomotion by mule deer and elk. *Journal of Wildlife Management*, 48(2):474–488, 1984. https://doi.org/10.2307/3801180.

Alex Potapov, Evelyn Merrill, Margo Pybus, and Mark A. Lewis. Chronic wasting disease: transmission mechanisms and the possibility of harvest management. *PLoS ONE*, 11(3):e0151039, 2016. https://doi.org/10.1371/journal.pone.0151039.

James Redcliffe, Rory Wilson, Mark Holton, Phil Hopkins, Mathieu Garel, Pascal Marchand, Gwendoline Wilson, Rowan Brown, and Luca Börger. Steep slopes, shallow angles: mountain ungulates create their own topography through movements. *Canadian Journal of Zoology*, 103:1–18, 2025. https://doi.org/10.1139/cjz-2024-0095.

Andrew Sharp and John Pastor. Stable limit cycles and the paradox of enrichment in a model of chronic wasting disease. *Ecological Applications*, 21(4):1024–1030, 2011. https://doi.org/10.1890/10-1449.1.

Arun Suresh, Zoe Shanley, Shane Hesting, Levi Jaster, Akila Raghavan, and Ram K. Raghavan. The direction and rate of spread of chronic wasting disease. *Frontiers in Veterinary Science*, 13: 1827624, 2026. https://doi.org/10.3389/fvets.2026.1827624.

James T. Tanner. The stability and the intrinsic growth rates of prey and predator populations. *Ecology*, 56(4):855–867, 1975. https://doi.org/10.2307/1936296.

Waldo Tobler. Three presentations on geographical analysis and modeling: Non-isotropic geographic modeling; speculations on the geometry of geography; and global spatial analysis. Technical Report 93-1, National Center for Geographic Information and Analysis, Santa Barbara, CA, 1993.

Fabienne D. Uehlinger, Anne C. Johnston, Trent K. Bollinger, and Cheryl L. Waldner. Systematic review of management strategies to control chronic wasting disease in wild deer populations in north america. *BMC Veterinary Research*, 12:173, 2016. https://doi.org/10.1186/s12917-016-0804-7.

U.S. Geological Survey. 1/3rd arc-second digital elevation models (DEMs) – USGS national map 3DEP downloadable data collection. U.S. Geological Survey, 2024a. URL https://www.sciencebase.gov/catalog/item/4f70aa9fe4b058caae3f8de5.

U.S. Geological Survey. Annual NLCD collection 1 science products (ver. 1.2, june 2026). U.S. Geological Survey data release, 2024b.

Gideon Wasserberg, Erik E. Osnas, Robert E. Rolley, and Michael D. Samuel. Host culling as an adaptive management tool for chronic wasting disease in white-tailed deer: a modelling study. *Journal of Applied Ecology*, 46(2):457–466, 2009. https://doi.org/10.1111/j.1365-2664.2008.01576.x.

Margaret A. Wild, N. Thompson Hobbs, Mark S. Graham, and Michael W. Miller. The role of predation in disease control: a comparison of selective and nonselective removal on prion disease dynamics in deer. *Journal of Wildlife Diseases*, 47(1):78–93, 2011. https://doi.org/10.7589/0090-3558-47.1.78.

Wei Yin, Wesley J. Marrero, Kamal Jnawali, Lale Asik, Michael G. Tyshenko, and Tamer Oraby. Spread of chronic wasting disease under stochastic environmental conditions and its control using deep reinforcement learning. *arXiv preprint arXiv:2609.05566*, 2026. https://doi.org/10.48550/arXiv.2609.05566.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- OnlineResource1.pdf