

ISBN 978-81-180544-2-9

WILDLIFE BIOLOGY AND CONSERVATION:

Science, Management and
Human-Wildlife Coexistence

DR. RANJIT R. RAUT
DR. KOMAL L. JAGZAP
DR. YOGESH V. BIDWE
MISS. KUMARI SHILPA PATEL

EDITORS



OLYMPICK
PUBLISHER INDIA

<https://olympickpublisher.in/>

SEPTEMBER
2026

Wildlife Biology and Conservation: Science, Management and Human- Wildlife Coexistence

Vol - I

2026

Dr. Ranjit R. Raut · Dr. Komal L. Jagzap

Dr. Yogesh V. Bidwe · Ms. Kumari Shilpa Patel



Editors

Dr. Ranjit R. Raut

Dr. Komal L. Jagzap

Dr. Yogesh V. Bidwe

Miss Kumari Shilpa Patel

DOI : <https://doi.org/10.5281/zenodo.22700647>

ISBN 978-81-180544-2-9

<https://olympickpublisher.in/>



contact@olympickpublisher.in

Registered Address: Olympick Publisher India 201, Shitole Empire Building,
New Sangvi, Haveli, Pune - 411027

© The Author(s) and Editor(s) (if applicable) hold an exclusive license from Olympick Publisher India. Any reproduction, distribution or transmission of this publication in any form requires prior written authorization from the publisher except when brief quotations are used in academic research or critical reviews.

At publication time the publisher, authors / editors made every possible effort to provide information and advice in this book which is both accurate and current. The authors/ editors give no expressed or implied warranties about the material's completeness or correctness and they accept no responsibility for mistakes or missing information.

When disposing of this product, please ensure proper paper recycling.

Contents

<i>Chapter – I</i> Foundations of Wildlife Ecology and Biology	<i>1-25</i>
<i>Chapter – II</i> Quantitative Wildlife Science and Monitoring	<i>26-51</i>
<i>Chapter – III</i> Habitat Ecology, Management and Restoration	<i>52-73</i>
<i>Chapter – IV</i> Protected Areas and Wildlife Resource Management	<i>74-94</i>
<i>Chapter – V</i> Human-Wildlife Conflict and Coexistence Frameworks	<i>95-116</i>
<i>Chapter – VI</i> Emerging Frontiers: Technology, Policy And Future Outlook	<i>117-137</i>

CHAPTER I: Foundations of Wildlife Ecology and Biology

Wildlife biology is the scientific study of free-living animal populations together with the ecological systems that sustain them. It draws on evolutionary theory, comparative physiology, demography, community ecology and landscape science. Its distinguishing feature is that it is a decision science as much as a descriptive one: the questions it asks are shaped by the need to predict how populations will respond to harvest, habitat alteration, disease, invasive species or a warming climate. This chapter establishes the conceptual and quantitative foundations on which the remainder of the book is built. Section 1.1 treats the individual organism as an evolved, physiologically constrained unit occupying a definable niche. Section 1.2 scales up to the population, the fundamental unit of management. Section 1.3 embeds populations in communities and landscapes. Section 1.4 places the resulting patterns in a global biogeographic and evolutionary context, then introduces the principal drivers of biodiversity loss along with the genetic principles that underpin recovery planning.




Figure 1.1. The hierarchy of organisation used throughout this chapter. Each level constrains the one to its right by supplying the mechanism, while each level to the right sets the context within which the mechanism operates. Section numbers indicate where each level is treated.

1.1 Principles of Wildlife Biology

1.1.1 Evolutionary adaptation

An adaptation is a heritable trait whose present form is the product of natural selection for the function it now performs. Selection acts on phenotypic variation that is heritable and that covaries with fitness, defined as the per capita contribution of a genotype to the next generation. Three consequences follow that matter directly for wildlife management.

First, adaptation is historically contingent. Populations converge not on globally optimal designs but on the best variant available in the local gene pool given the ancestral body plan, so the narrow gape of a specialist insectivore constrains the prey it can take even where larger

prey are abundant. Second, adaptation is local. Populations experiencing different selective regimes diverge genetically into locally adapted ecotypes; reciprocal transplant experiments repeatedly show a home-site advantage. This is the principal biological argument for sourcing translocation stock from environmentally matched populations, a point developed in Chapter 4.

Third, not all phenotypic variation is genetic. Phenotypic plasticity, the capacity of one genotype to express different phenotypes across environments, allows individuals to track short-term variation without genetic change. Plasticity buffers populations against climate variability, but the buffer has limits: where the rate of environmental change exceeds both the plastic and the evolutionary response capacity, decline follows.

Because energy acquired cannot be spent twice, allocation to growth, maintenance, storage and reproduction is subject to trade-offs. The resulting covariation among traits defines the fast-slow life-history continuum (Table 1.1). Fast species mature early, produce large litters or clutches, suffer high adult mortality and possess high intrinsic rates of increase. Slow species mature late, produce few offspring per attempt, survive well as adults and increase slowly. The position of a species on this continuum predicts a great deal about its management. Fast species tolerate substantial harvest yet fluctuate strongly with weather; slow species are demographically stable but recover from depletion over decades rather than seasons.

Table 1.1. The fast-slow life-history continuum, with indicative demographic values for representative vertebrates. Values are order-of-magnitude illustrations rather than population-specific estimates.

Attribute	Fast (e.g. field vole)	Intermediate (e.g. mallard)	Slow (e.g. African elephant)	Management implication
Age at first reproduction	5 to 6 weeks	1 year	11 to 14 years	Sets the lag before recovery begins
Litter or clutch size	4 to 8	9 to 13	1	Governs compensatory capacity
Adult annual survival	0.10 to 0.30	0.50 to 0.60	0.94 to 0.97	Determines elasticity structure
Intrinsic rate of increase r (per year)	3 to 5	0.4 to 0.8	0.02 to 0.06	Bounds sustainable harvest
Generation time	Under 0.5 year	2 to 3 years	25 to 30 years	Scales the timescale of monitoring
Principal source of variation	Environmental stochasticity	Mixed	Demographic stochasticity in small populations	Guides risk assessment

1.1.2 Physiological ecology

Physiological ecology explains why organisms occur where they do in terms of the mechanistic exchange of energy, water and heat with the environment. Its central variables are metabolic rate, thermal tolerance and water balance, all of which scale predictably with body size.

Basal metabolic rate B in endotherms scales with body mass M according to a power function with an exponent close to three quarters:

$$B = aM^{0.75}$$

Because the exponent is less than one, mass-specific metabolic rate declines with body size. A shrew must consume the equivalent of its own mass daily; an elephant may go days between meals of comparable relative size. This single relationship (Figure 1.2c) explains why small endotherms are vulnerable to short interruptions in food supply, why they cannot fast through severe weather without entering torpor and why minimum viable home range scales with mass.

The Scholander-Irving model describes how an endotherm meets its thermal costs (Figure 1.2a). Within the thermoneutral zone, bounded by the lower and upper critical temperatures, body temperature is maintained by adjusting insulation and posture at no metabolic cost above basal. Below the lower critical temperature, heat production must rise linearly with the temperature gradient. Above the upper critical temperature, evaporative cooling imposes a water cost that rises steeply. Winter range selection by ungulates, roost selection by bats and the timing of hibernation are all interpretable as behavioural strategies for remaining within or close to this zone.

Ectotherms are described instead by a thermal performance curve (Figure 1.2b), which is left-skewed and bounded by the critical thermal minimum and maximum. Performance rises gradually to an optimum, then collapses over a narrow interval above it. Two consequences are important for conservation under climate warming. Tropical ectotherms already operate close to their thermal optima, so even modest warming pushes them onto the descending limb; temperate ectotherms typically gain performance from the same warming. Thermal safety margins are therefore narrowest at low latitudes, an asymmetry that reverses the intuition that high-latitude species are always the most exposed.




Figure 1.2. Physiological constraints on distribution and abundance. (a) The Scholander-Irving model for an endotherm, showing the thermoneutral zone between the lower critical temperature and the upper critical temperature. (b) A thermal performance curve for an ectotherm, bounded by the critical thermal limits and peaking at the thermal optimum. (c) Metabolic allometry: basal metabolic rate scales with body mass raised to approximately 0.75, so mass-specific demand falls as body size rises.

1.1.3 The niche

The niche is the central organising concept linking an organism to its environment, but the word carries three distinct historical meanings that are still in active use (Table 1.2). Grinnell used it to mean the set of environmental conditions a species requires, a place-based concept. Elton used it to mean the functional role a species performs, principally its trophic position. Hutchinson unified these by defining the niche formally as an n -dimensional hypervolume, each axis being an environmental variable or resource, within which a population maintains a finite rate of increase of at least unity.

Table 1.2. Three niche concepts, their measurable expression and their principal use in wildlife science.

Concept	Definition	How it is measured	Principal application
Grinnellian	The set of abiotic conditions and habitat features a species requires	Occurrence records paired with climate and land-cover layers	Species distribution models; range-shift forecasting
Eltonian	The functional role of the species, principally its trophic position and impacts	Diet analysis; stable isotopes; interaction networks	Predicting community consequences of loss or introduction
Hutchinsonian	An n -dimensional hypervolume of conditions and resources permitting λ of at least 1	Multivariate ordination of measured niche axes	Formal analysis of overlap, breadth and coexistence

Hutchinson also drew the distinction that has proved most useful in practice. The fundamental niche is the hypervolume a species could occupy in the absence of interacting species. The realized niche is the smaller volume it actually occupies once competition, predation, parasitism and dispersal limitation are imposed (Figure 1.3a). The distinction matters because most correlative habitat models are fitted to occurrence data, so they estimate the realized niche. Extrapolating such a model to a site where the competitive context differs is hazardous, as is extrapolating it to a future climate in which that context will differ. Chapter 3 returns to this problem in the context of habitat suitability modelling.

The competitive exclusion principle states that two species with identical requirements cannot coexist indefinitely on a single limiting resource. Coexistence therefore implies niche partitioning along at least one axis: differences in prey size, foraging stratum, activity period or microhabitat. Where two species overlap broadly in part of their range but not elsewhere, selection may drive character displacement, an evolved divergence in the trait mediating competition. Niche breadth summarises how much of an axis a species exploits and separates generalists from specialists. Specialists are disproportionately represented among threatened taxa because a narrow niche offers no alternative when the resource it depends on is removed.




Figure 1.3. Niche concepts. (a) The fundamental niche is the n -dimensional hypervolume of conditions permitting positive population growth in the absence of interacting species. The realized niche is the subset actually occupied once competitors, predators and dispersal limitation are imposed. (b) Utilisation curves along a single resource axis, showing niche breadth, niche overlap and the partitioning that permits coexistence.

1.2 Population Dynamics

A population is a group of conspecific individuals occupying a defined area within which mating is effectively random. It is the operational unit of wildlife management because harvest quotas, recovery targets and viability assessments are all expressed in population terms. All change in population size N reduces to four processes, conventionally abbreviated BIDE:

$$N(t+1) = N(t) + B + I - D - E$$

where B is births, D is deaths, I is immigration and E is emigration. A closed population is one in which I and E are negligible, an assumption that simplifies estimation but is rarely satisfied over management-relevant timescales.

1.2.1 Life tables and vital rates

A life table organises age-specific survival and reproduction into a single accounting framework (Table 1.3). For each age class x it records the proportion of the original cohort still alive at the start of the class, denoted $l(x)$, together with the mean number of female offspring produced per surviving female, denoted $m(x)$. Three summary statistics follow directly.

The net reproductive rate is the mean number of female offspring produced by a female over her lifetime:

$$R_0 = \sum l(x) m(x)$$

A value of R_0 greater than one indicates that each female more than replaces herself. Generation time T , the mean age of the mothers of a cohort, is estimated as

$$T = \sum x l(x) m(x) / \sum l(x) m(x)$$

and the intrinsic rate of increase is approximated by $r = \ln(R_0) / T$. The finite rate of increase is $\lambda = \exp(r)$, so that λ greater than one denotes growth, λ equal to one denotes stability and λ less than one denotes decline. Reporting λ rather than r is standard

in wildlife science because it is directly interpretable as a multiplier of population size per time step.

Table 1.3. A cohort life table for a hypothetical large herbivore, based on 1000 female calves. Survivorship $l(x)$ is the proportion of the original cohort alive at the start of age class x ; $m(x)$ is female offspring per female. Summing the $l(x)m(x)$ column gives $R_0 = 1.77$. The age-weighted sum of 10.92 divided by R_0 yields a generation time of 6.2 years, so $r = 0.092$ and $\lambda = 1.096$.

Age (years)	x	Number alive	Survivorship $l(x)$	Fecundity $m(x)$	$l(x) m(x)$	$x l(x) m(x)$
0		1000	1.000	0.00	0.000	0.00
1		620	0.620	0.00	0.000	0.00
2		540	0.540	0.10	0.054	0.11
3		500	0.500	0.40	0.200	0.60
4		470	0.470	0.55	0.259	1.04
5		440	0.440	0.60	0.264	1.32
6		410	0.410	0.60	0.246	1.48
7		375	0.375	0.60	0.225	1.58
8		335	0.335	0.55	0.184	1.47
9		285	0.285	0.50	0.143	1.29
10		225	0.225	0.45	0.101	1.01
11		155	0.155	0.40	0.062	0.68
12		80	0.080	0.30	0.024	0.29
13		25	0.025	0.20	0.005	0.07
Total					$R_0 = 1.767$	$\Sigma = 10.92$

1.2.2 Exponential and logistic growth

In a constant environment with unlimited resources, growth is exponential:

$$dN/dt = rN, \text{ so that } N(t) = N(0) e^{(rt)}$$

Exponential growth is not a biological curiosity. Reintroduced populations, invading species and populations released from a limiting predator all display it over the initial phase of establishment, sometimes for decades. It cannot persist indefinitely because resources are finite. Introducing a linear reduction in per capita growth with density yields the logistic model:

$$dN/dt = rN (1 - N/K)$$

where K is the carrying capacity. The trajectory is sigmoidal, with maximum absolute recruitment at $N = K/2$ (Figure 1.4a). This inflexion point is the theoretical basis of maximum sustainable yield, treated in Chapter 4.

The assumption of linear density dependence embedded in the logistic model is seldom met. The theta-logistic generalisation relaxes it:

$$dN/dt = rN [1 - (N/K)^\theta]$$

Values of θ below one produce strong density dependence at low abundance, characteristic of fast-lived species; values above one concentrate the density response near K , a pattern typical of large, slow-lived vertebrates (Figure 1.4b). Fitting the exponent rather than assuming it is now standard practice, because the shape of the density response determines the abundance at which maximum yield occurs.




Figure 1.4. Models of population growth. (a) Exponential growth compared with logistic growth toward a carrying capacity K , with maximum absolute recruitment at $K/2$. (b) The theta-logistic family: theta below one shifts density dependence toward low abundance, while theta above one concentrates it near K . (c) A strong Allee effect introduces a critical threshold A below which the per capita growth rate is negative, so that decline becomes deterministic.

1.2.3 Carrying capacity and density dependence

Carrying capacity is used in at least three distinct senses in wildlife science, which are frequently conflated. Ecological carrying capacity is the equilibrium abundance a habitat will support in the absence of harvest. Economic or maximum sustained yield carrying capacity is the abundance that maximises sustainable offtake, typically near $K/2$. Cultural carrying capacity is the abundance a human community will tolerate, which for species that damage crops or transmit disease may be far below the ecological value. Conflict management, treated in Chapter 5, is largely the practice of reconciling the third with the first.

Density dependence is the mechanism that generates carrying capacity. As density rises, per capita food declines, competition for breeding sites intensifies, parasite transmission increases and social stress rises. Eberhardt observed that in large mammals these responses appear in a predictable sequence as density approaches K : juvenile survival falls first, then age at first reproduction increases, then adult reproductive rate declines and only finally does adult survival decrease. This ordering is diagnostic and is used to infer proximity to carrying capacity from demographic data alone.

The distinction between compensatory and additive mortality follows directly from density dependence. Under compensation, mortality from one source, such as regulated harvest, is offset by reduced mortality from other sources because the survivors experience relaxed competition. Under addition, harvest mortality is superimposed on natural mortality with no

compensating response. Real populations lie between these extremes, with compensation strongest well below K . Table 1.4 summarises the practical distinctions.

Table 1.4. Forms of density dependence and their consequences for management.

Process	Description	Consequence for management
Negative density dependence	Per capita growth declines as density rises through resource competition, disease or interference	Generates a stable equilibrium; permits sustainable harvest
Compensatory mortality	Harvest mortality is offset by reduced natural mortality among survivors	Moderate harvest may not reduce autumn-to-autumn abundance
Additive mortality	Harvest mortality adds directly to natural mortality	Every animal removed reduces the subsequent population
Density independence	Weather, fire or flood affects vital rates irrespective of abundance	Produces fluctuation without regulation; raises extinction risk
Allee effect (positive density dependence)	Per capita growth declines as density falls, through mate limitation, failed group defence or reduced foraging efficiency	Creates a threshold below which recovery is impossible without intervention

The Allee effect deserves separate emphasis because it inverts the usual logic. Below a critical density, individuals may fail to locate mates, cooperative hunting or communal defence may break down; reproductive synchrony may also collapse. A strong Allee effect produces a threshold density A below which the per capita growth rate is negative regardless of resource availability (Figure 1.4c). The passenger pigeon, a colonial breeder dependent on enormous aggregations for predator satiation, is the canonical example. Recognising an Allee threshold changes the objective of recovery from slowing decline to lifting the population above the threshold in a single decisive intervention.

1.2.4 Structured models

Because vital rates differ among age or stage classes, unstructured models discard information that managers need. Matrix projection models retain it. The population vector $n(t)$ records abundance in each class, while a projection matrix A holds survival probabilities on the subdiagonal together with fecundities in the top row. Projection proceeds as $n(t+1) = A n(t)$. The dominant eigenvalue of A is λ , the asymptotic growth rate, while the corresponding right eigenvector gives the stable stage distribution.

The practical value of the matrix approach lies in perturbation analysis. Elasticity measures the proportional change in λ resulting from a proportional change in a matrix element, so elasticities are directly comparable across vital rates. A robust generalisation emerges: in long-lived, slow-reproducing species, λ is most elastic to adult survival, whereas in short-lived,

fast-reproducing species it is most elastic to fecundity and juvenile survival. This explains why sea turtle conservation shifted from egg protection to reducing adult bycatch, why raptor recovery hinges on adult mortality from poisoning or collision and why waterfowl management can rely on regulated harvest of a fecund resource.

1.2.5 Metapopulation dynamics

Few wildlife populations occupy continuous habitat. Where suitable habitat occurs as discrete patches connected by dispersal, the appropriate unit of analysis becomes the metapopulation: a set of local populations linked by movement, each subject to local extinction with recolonisation from occupied patches. Levins expressed the classical case as a balance between colonisation and extinction of patches:

$$dp/dt = c p (1 - p) - e p$$

where p is the proportion of patches occupied, c is the colonisation rate and e the local extinction rate. The equilibrium occupancy is $p^* = 1 - e/c$. Two results follow that are central to reserve design. Regional persistence requires c to exceed e , so a metapopulation can persist indefinitely even though every local population is ephemeral. Equally, persistence never requires all patches to be occupied: an equilibrium occupancy of 0.4 is a functioning metapopulation, not a failing one, so unoccupied patches retain conservation value as recolonisation targets.




Figure 1.5. Metapopulation structure and dynamics. (a) The classical Levins structure, with similar patches subject to local extinction and recolonisation. (b) Mainland-island structure, in which a large persistent population supplies satellites. (c) Source-sink structure, in which sinks with λ below one persist only through immigration from sources. (d) Levins dynamics for three combinations of colonisation and extinction rate, with equilibrium occupancy marked. (e) Relaxation after habitat loss, illustrating extinction debt: species loss lags the loss of habitat, so the shaded area represents extinctions that are committed but not yet realised.

Real systems depart from the classical case in ways summarised in Table 1.5. In mainland-island systems a large persistent population supplies satellites, so regional dynamics are dominated by the mainland. In source-sink systems, patches differ in quality: sources have

lambda greater than one and export surplus individuals, whereas sinks have lambda less than one and persist only through immigration. The distinction is critical for monitoring, because a sink can support a substantial and apparently stable population while contributing nothing to regional persistence. Protecting an occupied sink while the source is lost guarantees eventual collapse.

Table 1.5. Metapopulation structures, their diagnostic features and their management implications.

Structure	Diagnostic features	Dominant process	Management priority
Classical (Levins)	Patches of similar size and quality; frequent turnover	Balance of colonisation against extinction	Maintain patch number and connectivity
Mainland-island	One large patch with persistent occupancy plus small satellites	Unidirectional dispersal from the mainland	Protect the mainland patch above all
Source-sink	Patches differ in intrinsic quality; occupancy is not evidence of quality	Net dispersal from sources to sinks	Identify and secure sources; avoid ecological traps
Patchy	Dispersal so frequent that patches function as one population	Movement swamps local dynamics	Manage as a single unit
Non-equilibrium	Extinction consistently exceeds colonisation	Regional decline toward extirpation	Restore connectivity or accept translocation

Two further phenomena follow from metapopulation structure. The rescue effect describes immigration into an occupied patch that lowers its extinction probability, so that connectivity raises persistence even without recolonisation. Extinction debt describes the lag between habitat loss and the eventual loss of species that habitat could no longer support (Figure 1.5e). Because relaxation may take decades in long-lived taxa, present-day species lists from recently fragmented landscapes systematically overstate long-term viability. Conversely, the same lag creates an opportunity: restoration undertaken before the debt is paid can prevent extinctions that are already committed.

1.3 Community and Landscape Ecology

1.3.1 Species interactions

A community is an assemblage of populations that co-occur and interact. Interactions are conventionally classified by the sign of the net effect on each participant (Table 1.6). This sign convention is coarse, but it disciplines reasoning about indirect effects, which are frequently

larger than direct ones. Removing a predator has a direct positive effect on its prey plus an indirect negative effect on the plants that prey consumes, an indirect positive effect on competitors of that prey and an indirect effect of uncertain sign on smaller predators formerly suppressed by the apex species.

Table 1.6. The sign convention for pairwise species interactions, with wildlife examples.

Interaction	Signs	Mechanism	Wildlife example
Predation, parasitism, herbivory	+ / –	One participant consumes part or all of the other	Wolf and elk; tick and moose
Competition	– / –	Shared limiting resource, exploitative or interference	Coyote and swift fox for prey and dens
Mutualism	+ / +	Reciprocal benefit	Frugivorous hornbill and fruiting tree
Commensalism	+ / 0	One benefits; the other is unaffected	Cattle egret following grazing ungulates
Amensalism	– / 0	One is harmed; the other is unaffected	Trampling of nests by large herbivores
Intraguild predation	+ / –	A predator kills a competing predator	Wolf killing coyote

1.3.2 Predator and prey

The consumption rate of an individual predator as a function of prey density is its functional response (Figure 1.6a). Three forms are recognised. The Type I response rises linearly to a ceiling and is confined to filter feeders. The Type II response, described by the disc equation, rises with decelerating slope to an asymptote set by handling time:

$$f(N) = aN / (1 + a h N)$$

where a is the attack rate and h the handling time per prey item. Because the proportion of prey killed declines monotonically with prey density, a Type II response is destabilising: predation pressure is heaviest precisely when prey are scarce. The Type III response is sigmoidal, arising from prey switching, formation of a search image or the presence of refuges at low prey density. Because the proportion killed rises with density over the lower range, a Type III response can regulate prey.

The numerical response describes change in predator abundance with prey density, through reproduction or aggregation. Coupling functional and numerical responses generates the delayed density dependence responsible for cyclic dynamics such as the ten-year lynx and snowshoe hare cycle of the boreal forest (Figure 1.6b). Where a generalist predator maintains high abundance on alternative prey while imposing a Type III response on a declining secondary

prey species, the outcome may be a predator pit: a low-density stable equilibrium from which the prey population cannot escape without intervention (Figure 1.6c). Predator pits explain why simply removing the original cause of a decline sometimes fails to produce recovery.




Figure 1.6. Predation. (a) The three functional responses. The Type II form is destabilising because the proportion of prey taken falls as prey density rises; the Type III form can regulate prey over the lower part of its range. (b) Coupled predator-prey cycles generated by a Lotka-Volterra system. (c) A predator pit: prey recruitment intersects a Type III predation curve three times, giving a low stable equilibrium, an unstable threshold and a high stable equilibrium near K .

1.3.3 Trophic cascades

A trophic cascade occurs when the effect of a consumer propagates through two or more trophic levels. Top-down control means that consumers limit the level below; bottom-up control means that resource supply limits the level above. Both operate simultaneously in most systems, with their relative strength varying by productivity and system openness.

Keystone species exert effects on community structure disproportionate to their biomass. Ecosystem engineers such as beaver create habitat for other species through physical modification. The removal of large predators, termed trophic downgrading, is the most pervasive human alteration of terrestrial communities. Its consequences include herbivore release with subsequent suppression of woody recruitment, along with mesopredator release, in which medium-sized predators formerly limited by intraguild predation increase and impose heavy predation on small vertebrates.

The reintroduction of wolves to the Greater Yellowstone Ecosystem from 1995 provides the most examined terrestrial case (Figure 1.7). Wolves reduced elk numbers while also changing elk distribution and vigilance, which relaxed browsing pressure in riparian zones. Willow and aspen recruitment increased in places, supporting beaver colonisation, which in turn raised water tables and created habitat for riparian birds, amphibians and fish. Wolves also suppressed coyotes, which is associated with improved pronghorn fawn survival. Two cautions are essential: the magnitude of the woody-plant response has been contested, since drought, bison grazing and climate all contribute; and the cascade is context dependent, so it should not be assumed to transfer elsewhere without evidence.




Figure 1.7. The Yellowstone trophic cascade, contrasting the state of the system with the apex predator absent (left) against its state following restoration (right). Arrows indicate the direction of propagation; the symbols indicate whether the receiving variable increases or decreases. The lateral arrows show intraguild predation, whose loss releases the mesopredator.

1.3.4 Landscape structure and scale

Landscape ecology examines how spatial pattern influences ecological process. Its standard abstraction is the patch-corridor-matrix model: discrete patches of habitat, linear corridors that connect them and a surrounding matrix of unsuitable or less suitable cover. The matrix is not inert. Its permeability determines whether corridors function, so a hedgerow crossing intensive arable land performs differently from the same hedgerow crossing rough pasture.

Two scale parameters must be specified in any landscape analysis. Grain is the finest resolution at which data are recorded; extent is the total area encompassed. A pattern detectable at one grain may be invisible at another, so the scale of measurement should match the scale at which the focal organism perceives the landscape, usually approximated from home range size or dispersal distance. Table 1.7 lists the metrics most often used to summarise landscape structure.

Table 1.7. Common landscape metrics, their definitions and their ecological interpretation.

Metric	Definition	Ecological interpretation
Percentage habitat cover	Proportion of the landscape in the focal class	Primary predictor of occupancy; thresholds are common near 20 to 30 per cent
Patch density	Number of patches per unit area	Rises then falls as fragmentation proceeds to attrition
Mean patch size	Average area of habitat patches	Predicts presence of area-sensitive species
Edge density	Length of habitat boundary per unit area	Indexes exposure to edge-mediated processes
Core area	Patch area remaining after removing the edge zone	Better predictor than total area for interior specialists
Connectivity or patch cohesion	Degree to which patches are functionally joined for a given disperser	Governs recolonisation and gene flow

1.3.5 Edge effects

An edge is the boundary between two cover types. Edge effects are the biotic and abiotic changes that occur near that boundary. They penetrate the patch to a characteristic distance termed the depth of edge influence (Figure 1.8). Three classes are distinguished. Abiotic edge effects include elevated solar radiation, higher air and soil temperature, lower humidity and greater wind exposure; these typically penetrate tens of metres. Direct biotic effects include the replacement of shade-tolerant interior vegetation by light-demanding species.

Depth of edge influence varies by more than an order of magnitude among responses (Table 1.8), so no single buffer width protects against all of them. The operational consequence is the core area concept: for a specified depth d , the core is the patch area remaining after a buffer of width d is removed from the perimeter. Because core area scales with the square of linear dimension while edge scales linearly, small patches may contain no core at all. A circular patch of 10 hectares subject to a 100 metre edge zone retains roughly 1.9 hectares of core, whereas the same edge zone applied to a 1 hectare patch removes the patch entirely.




Figure 1.8. Edge effects. (a) The core area of a patch is what remains once a buffer of width equal to the depth of edge influence is removed. Because edge scales with the perimeter while core scales with area, small patches may retain no core at all. (b) Different responses decay at different rates, so the depth of edge influence is response specific rather than a single property of the boundary.

Table 1.8. Indicative depth of edge influence for different classes of response in temperate forest. Values summarise the range commonly reported and should be verified locally.

Response	Typical depth	Notes
Wind throw and canopy damage	10 to 60 m	Greatest on recently created hard edges
Air temperature, humidity, light	15 to 50 m	Declines as edge vegetation seals the boundary
Soil moisture and litter depth	10 to 30 m	Interacts strongly with aspect
Invasive plant establishment	20 to 100 m	Tracks propagule pressure from the matrix
Nest predation and brood parasitism	50 to 300 m	Highest where the matrix supports generalist predators
Avoidance by interior forest birds	100 to 400 m	Behavioural, so it may exceed the abiotic depth

1.3.6 Habitat fragmentation

Fragmentation is the process by which continuous habitat is broken into smaller, more isolated remnants. It proceeds through recognisable stages: perforation, dissection, subdivision and attrition (Figure 1.9a to 1.9d). It is essential to separate habitat loss from fragmentation per se, meaning the breaking apart of habitat at a constant total area. Meta-analyses consistently find that habitat loss is the dominant negative effect, whereas the additional effect of fragmentation per se is weaker and inconsistent in sign.

The theoretical foundation for analysing fragmented systems is the equilibrium theory of island biogeography, which predicts species richness on an island as a balance between two rates: immigration, which declines with distance from the mainland; and extinction, which declines with island area (Figure 1.10a). The empirical expression of the same relationship is the species-area curve:

$$S = c A^z$$

The exponent z is typically about 0.15 for nested samples within a continuous area, about 0.25 for habitat patches and about 0.35 for true islands (Figure 1.10b). Higher z means richness falls faster as area is lost, which is why habitat remnants surrounded by a hostile matrix behave more like islands than like samples of continuous habitat.

The theory generated the SLOSS debate, on whether a single large reserve or several small reserves of equal total area conserves more species. There is no universal answer. A single large reserve better supports area-sensitive species, wide-ranging predators plus viable populations of low-density taxa; several small reserves may capture greater habitat heterogeneity, spread risk against catastrophe and conserve more species overall where beta diversity among sites is high. Contemporary reserve design accordingly treats the question empirically for the taxa and landscape in hand, an approach developed in Chapter 4.




Figure 1.9. Fragmentation. (a) to (d) The stages of fragmentation, from an intact block through perforation and dissection to attrition, with the proportion of habitat remaining. (e) The cascade of demographic, genetic and community consequences that follows the loss and subdivision of habitat.




Figure 1.10. Island biogeography. (a) The equilibrium model. Immigration falls as more species accumulate and is lower for distant islands; extinction rises with richness and is higher on small islands. Equilibrium richness occurs where the curves intersect, giving four combinations of area and isolation. (b) The species-area relationship on logarithmic axes, where the slope is the exponent z .

1.4 Biodiversity Patterns and Threats

1.4.1 Dimensions and measurement of biodiversity

Biodiversity is the variety of life at three nested levels: genetic diversity within populations, species diversity within communities and ecosystem diversity within landscapes and regions. Conservation planning that addresses only the species level is incomplete, because genetic diversity determines adaptive potential while ecosystem diversity determines the persistence of the processes that generate both.

Species diversity is partitioned into three components. Alpha diversity is richness within a site. Beta diversity is turnover in composition among sites. Gamma diversity is total richness across a region. The partition matters operationally: a region with high beta diversity requires a distributed network of reserves, because no single site can represent the regional biota. Richness alone ignores relative abundance, so evenness-weighted indices are also used. The Shannon index $H = -\sum p(i) \ln p(i)$ is sensitive to rare species, whereas the Simpson index $D = 1 / \sum p(i)^2$ is dominated by common species. Reporting both, alongside raw richness, is good practice.

1.4.2 Biogeographic patterns

The most general pattern in biology is the latitudinal diversity gradient: species richness in most taxa rises from the poles toward the equator (Figure 1.11a). Numerous hypotheses have been proposed to explain it. The species-energy hypothesis holds that greater available energy supports more individuals and therefore more species. The evolutionary time and area hypothesis holds that tropical biomes are older and have been larger, allowing more speciation with less extinction. The climatic stability hypothesis holds that lower seasonality permits narrower niches. The biotic interactions hypothesis holds that stronger predation and competition promote specialisation. These explanations are not mutually exclusive; current

evidence supports a combination in which energy availability and evolutionary history are the leading terms.

Related regularities are widely used in planning. Rapoport's rule states that range sizes tend to increase with latitude, so tropical species are typically narrow-ranged and hence more extinction prone. Elevational gradients frequently show a mid-elevation peak in richness rather than a monotonic decline (Figure 1.11b), a pattern attributed to the intersection of area, energy and geometric constraints. Endemism is concentrated in isolated systems such as oceanic islands, ancient lakes and montane sky islands. The combination of high endemism with high habitat loss defines the biodiversity hotspot concept, which has been used to prioritise a small proportion of the terrestrial surface holding a large share of endemic species.

Biogeographic realms remind us that current distributions record deep history, expressed most sharply in faunal discontinuities such as Wallace's Line separating the Asian from the Australasian fauna. Two ecologically similar sites in different realms will support unrelated faunas, which is why conservation value cannot be judged from species richness alone and must incorporate irreplaceability, the extent to which a site holds features found nowhere else.





Figure 1.11. Global patterns of diversity and threat. (a) The latitudinal diversity gradient. (b) Two commonly observed elevational patterns. (c) The proportion of assessed species classified as threatened, that is, Critically Endangered, Endangered or Vulnerable, for selected groups. Source: IUCN Red List, 2025 assessment summary, based on more than 169,000 species assessed of which more than 47,000 are threatened.

1.4.3 Extinction: rates, process and assessment

Extinction is a normal evolutionary outcome. The background rate, estimated from the fossil record, is conventionally expressed as extinctions per million species-years. Estimates of the background vertebrate rate cluster near 0.1 to 2 extinctions per million species-years, whereas documented modern rates for well-monitored groups such as birds, mammals and amphibians exceed this by one to three orders of magnitude. The pattern of loss also differs from the background pattern: it is concentrated in large-bodied, slow-reproducing, narrow-ranged and specialised taxa, which is the signature of a driver acting through habitat and exploitation rather than through random environmental change.

The process by which a small population is driven to extinction is the extinction vortex (Figure 1.12). A reduction in population size raises the intensity of genetic drift plus inbreeding, which depresses survival and fecundity, which lowers λ , which reduces population size further

while amplifying the influence of demographic and environmental stochasticity. Because each turn of the cycle tightens the next, the process is self-reinforcing. Recognising a population as being in the vortex changes management from habitat provision alone to direct demographic and genetic intervention.



Each turn of the cycle reduces population size further, so the process is self-reinforcing.

Figure 1.12. The extinction vortex. Genetic and demographic processes reinforce one another in small populations, so each turn of the cycle reduces population size further. Intervention is most effective at the top of the cycle, before genetic effects become the limiting factor.

Extinction risk is assessed formally through the IUCN Red List, which applies five quantitative criteria to place taxa in categories of increasing risk (Table 1.9). The system is deliberately criterion-based rather than expert-opinion based, so that assessments are repeatable and comparable across taxa. A taxon qualifies for a category if it meets any one criterion, which allows assessment where only some data are available. Two cautions apply. Categories describe relative extinction risk rather than conservation priority, which must also weigh cost, tractability and distinctiveness. And the Data Deficient category is not a statement of low risk; Data Deficient taxa are, on average, at least as threatened as assessed ones.

Table 1.9. The IUCN Red List categories and a simplified summary of the quantitative thresholds for the threatened categories. Full criteria and qualifiers are given in the current IUCN guidelines.

Category	Code	Criterion A: population reduction	Criterion C or D: small population
Critically Endangered	CR	At least 80 per cent over 10 years or three generations	Fewer than 250 mature individuals; fewer than 50 under D
Endangered	EN	At least 50 per cent over the same period	Fewer than 2500 mature individuals; fewer than 250 under D
Vulnerable	VU	At least 30 per cent over the same period	Fewer than 10,000 mature individuals; fewer than 1000 under D
Near Threatened	NT	Close to qualifying for Vulnerable	Close to qualifying for Vulnerable
Least Concern	LC	Does not qualify under any criterion	Does not qualify under any criterion
Data Deficient	DD	Information inadequate for assessment	Not a category of low risk
Extinct in the Wild / Extinct	EW / EX	Surviving only in cultivation or captivity; or no reasonable doubt the last individual has died	


1.4.4 The HIPPO framework

The proximate drivers of biodiversity loss are conveniently summarised by the mnemonic HIPPO: Habitat destruction, Invasive species, Pollution, human Population and Overharvesting (Figure 1.13, Table 1.10). The ordering is approximately that of global importance for terrestrial vertebrates, although the ranking differs by realm and by taxon. Habitat destruction is the leading driver for most terrestrial and freshwater taxa; overharvesting leads for many marine species along with large-bodied tropical vertebrates; invasive species dominate on oceanic islands, where they are implicated in the majority of recorded vertebrate extinctions.

Climate change is not a separate letter in the mnemonic, which reflects the era in which the framework was devised rather than its current importance. It is best treated as a threat multiplier that interacts with each of the five drivers: shifting ranges collide with fragmented habitat, warming assists invasive species, altered hydrology concentrates pollutants and phenological mismatch reduces recruitment. Chapter 3 treats adaptation responses in detail.

Table 1.10. The HIPPO drivers, their principal mechanisms and representative wildlife examples.

Driver	Principal mechanisms	Representative examples
Habitat destruction, degradation and fragmentation	Conversion to agriculture, logging, dams, urban expansion, road networks	Orangutans and lowland rainforest conversion; grassland birds and arable intensification
Invasive alien species	Predation, competition, disease transmission, hybridisation, habitat modification	Brown tree snake in Guam; avian malaria in Hawaii; feral cats and rodents on islands
Pollution	Nutrient enrichment, pesticides, heavy metals, plastics, light and noise	Diclofenac and Asian vulture collapse; neonicotinoids and insectivore declines
Human population and consumption	Infrastructure, resource demand, disturbance, expanding interface with wildlife	Road mortality; recreational disturbance of breeding shorebirds
Overharvesting	Hunting, fishing, wildlife trade, incidental take	Pangolin trade; elephant and rhinoceros poaching; albatross bycatch in longline fisheries

**Figure 1.13.** The HIPPO framework. Human population growth together with rising per capita consumption drives five proximate mechanisms of biodiversity loss, which converge on reduced abundance, contracted range and fragmented populations. Climate change is shown as a cross-cutting multiplier acting on all five.

1.4.5 Conservation genetics baseline

Genetic diversity is the raw material of adaptation, so its retention is a legitimate conservation objective in its own right. It is usually summarised by expected heterozygosity, the proportion

of loci at which a randomly chosen individual would be heterozygous under Hardy-Weinberg expectations, together with allelic richness, which is more sensitive to recent bottlenecks because rare alleles are lost before heterozygosity declines measurably. Table 1.11 defines the core parameters.

Table 1.11. Core parameters of conservation genetics.

Parameter	Definition	Why it matters
Expected heterozygosity, H_e	Proportion of loci expected to be heterozygous under Hardy-Weinberg equilibrium	Standard index of standing genetic variation
Allelic richness	Number of alleles per locus, corrected for sample size	Detects recent bottlenecks earlier than H_e does
Effective population size, N_e	Size of an ideal population losing variation at the same rate as the observed one	Governs drift and inbreeding; usually far below census size
Inbreeding coefficient, F	Probability that two alleles at a locus are identical by descent	Predicts the magnitude of inbreeding depression
Gene flow, N_m	Number of effective migrants exchanged per generation	One to ten migrants per generation limits differentiation
Genetic load	Accumulated deleterious recessive alleles in a population	Expressed as inbreeding depression when homozygosity rises

Genetic drift is the random change in allele frequency between generations arising from finite sampling of gametes. Its rate is governed not by census size N but by effective population size N_e . Under drift alone, heterozygosity declines each generation by a factor of $1/(2N_e)$:

$$H(t) = H(0) [1 - 1/(2N_e)]^t$$

The consequences are severe at small N_e (Figure 1.14a). A population with N_e of 500 retains about 90 per cent of its heterozygosity after 100 generations, whereas a population with N_e of 50 retains only about 37 per cent. Effective size is almost always smaller than census size, typically by an order of magnitude, because of unequal sex ratio, variance in reproductive success, fluctuation in abundance across generations and overlapping generations. For unequal sex ratio the relationship is

$$N_e = 4 N_m N_f / (N_m + N_f)$$

so a polygynous population of 200 animals in which only 5 males breed has an effective size near 20, not 200 (Figure 1.14c). Where abundance fluctuates, the long-run effective size is the harmonic mean across generations, which is dominated by the smallest values. A single severe bottleneck therefore leaves a genetic signature that persists long after demographic recovery

(Figure 1.14b), which is why the northern elephant seal retains very low variation despite a present population in the hundreds of thousands.




Figure 1.14. Genetic processes in small populations. (a) Loss of heterozygosity under drift for four values of effective population size. (b) A bottleneck followed by demographic recovery: census size returns, but heterozygosity does not. (c) The effect of unequal sex ratio on effective size in a population of 200 individuals.

Inbreeding, the mating of relatives, raises homozygosity, which exposes deleterious recessive alleles and reduces heterozygote advantage. The resulting decline in survival, fecundity or disease resistance is inbreeding depression. It has been documented across wild vertebrates and can be substantial: the Isle Royale wolf population, isolated for decades, developed skeletal deformities together with severely reduced recruitment before collapsing. The converse risk, outbreeding depression, arises when crossing genetically divergent populations breaks up locally adapted gene complexes. It is real but far rarer than inbreeding depression. Current guidance is that the risk of inaction usually exceeds the risk of managed gene flow between populations that have been separated for fewer than about 500 years with no fixed chromosomal differences.

The classical rule of thumb, the 50/500 rule, proposed that an effective size of 50 is needed to avoid immediate inbreeding depression while 500 is needed to retain evolutionary potential. Both figures have been revised upward, with 100/1000 now more widely advocated. Both should be treated as indicative rather than as targets to be defended. The operational lesson is unchanged: because N_e is typically about one tenth of census size, a population must usually number in the thousands to be genetically secure over evolutionary timescales.

Genetic rescue is the deliberate introduction of individuals from a related population to restore variation and relieve inbreeding depression. The Florida panther programme, in which eight female pumas from Texas were released into a population of roughly 25 highly inbred animals in 1995, produced measurable increases in heterozygosity, a decline in the frequency of congenital defects and a substantial increase in population size. Genetic rescue is now an accepted tool for small isolated populations, subject to careful assessment of source-population match. Chapter 4 develops the operational protocols.

Chapter Summary

- Adaptation is heritable, local and historically contingent, so populations are constrained by their evolutionary history as well as by present conditions. Life-history trade-offs place species on a fast-slow continuum that predicts both harvest tolerance and recovery rate.
- Metabolic rate scales with body mass raised to about 0.75, while thermal tolerance is bounded by critical limits. These constraints set the coarse limits of distribution. They also explain why tropical ectotherms carry the narrowest thermal safety margins under warming.
- The realized niche is smaller than the fundamental niche because of competition, predation and dispersal limitation. Correlative habitat models estimate the realized niche, which limits their transferability.
- Population change reduces to births, deaths, immigration and emigration. Life tables yield R_0 , generation time and λ ; matrix models add stage structure plus elasticity, which identifies the vital rate on which management should act.
- Density dependence generates carrying capacity and permits sustainable harvest through compensation, while positive density dependence at low abundance produces Allee thresholds that make decline self-accelerating.
- Metapopulation theory shows that regional persistence can coexist with routine local extinction, provided colonisation exceeds extinction. Source-sink structure means that occupancy is not evidence of habitat quality.
- Functional and numerical responses determine whether predation regulates or destabilises prey. Trophic cascades and mesopredator release show that removing apex predators restructures whole communities.
- Habitat loss is the dominant fragmentation effect, but configuration governs edge exposure, core area and connectivity. Core area, not total area, predicts the persistence of interior specialists.
- Species richness peaks in the tropics for reasons combining energy availability with evolutionary history, while current extinction rates exceed background rates by one to three orders of magnitude and are biased toward large, slow, narrow-ranged taxa.
- Genetic drift proceeds at a rate set by effective population size, typically about one tenth of census size, so small populations lose variation, express inbreeding depression and enter a self-reinforcing extinction vortex unless managed.

Key Terms

Allee effect positive density dependence at low abundance, producing a threshold below which growth is negative

Carrying capacity (K) the equilibrium abundance a habitat supports; distinguish ecological from cultural carrying capacity

Core area patch area remaining after the edge zone is removed

Depth of edge influence the distance to which edge-mediated change penetrates a patch

Effective population size (N_e) the size of an ideal population losing genetic variation at the observed rate

Extinction debt the lag between habitat loss and the eventual loss of species it can no longer support

Extinction vortex the self-reinforcing interaction of small size, genetic erosion and stochasticity

Functional response the relationship between prey density and per capita consumption by a predator

Fundamental and realized niche the conditions permitting positive growth without interacting species; the portion of that volume actually occupied

HIPPO habitat destruction, invasive species, pollution, human population and overharvesting

Inbreeding depression reduced fitness arising from increased homozygosity

Metapopulation a set of local populations linked by dispersal, with local extinction and recolonisation

Source and sink patches with λ above one that export individuals; patches with λ below one that depend on immigration

Trophic cascade the propagation of a consumer effect through two or more trophic levels

Review Questions

1. Using the life table in Table 1.3, recompute R_0 and λ under a scenario in which first-year survival falls from 0.62 to 0.45. State in one sentence what this implies for the population trajectory.
2. A correlative habitat model fitted to occurrence data performs poorly when transferred to a neighbouring region. Explain this outcome in terms of the distinction between the fundamental and the realized niche.
3. Explain why a Type II functional response is destabilising while a Type III response can regulate prey, then state the circumstances under which a generalist predator produces a predator pit.
4. A 25 hectare circular woodland is surrounded by intensive arable land. Assuming a depth of edge influence of 150 metres, calculate the core area. Comment on the implications for an interior-nesting songbird.
5. Distinguish habitat loss from fragmentation per se. Design a study that would separate their effects on a bird community, stating your response variables plus the landscape metrics you would measure.
6. A managed ungulate population of 400 animals has an operational sex ratio of 15 breeding males to 185 breeding females. Estimate N_e and comment on the rate at which heterozygosity will be lost.

7. Choose a species you know well. Place it on the fast-slow continuum, identify the vital rate to which λ is most elastic and state one management action that follows.

Further Reading

- Begon, M. and Townsend, C. R. Ecology: From Individuals to Ecosystems. Comprehensive treatment of the theory in Sections 1.2 and 1.3.
- Caswell, H. Matrix Population Models: Construction, Analysis and Interpretation. The standard reference for structured models plus perturbation analysis.
- Frankham, R., Ballou, J. D. and Briscoe, D. A. Introduction to Conservation Genetics. Full development of the material in Section 1.4.5, including effective size estimation and genetic rescue.
- Fahrig, L. Effects of habitat fragmentation on biodiversity. Annual Review of Ecology, Evolution and Systematics. The key synthesis separating habitat loss from fragmentation per se.
- IUCN Red List of Threatened Species, current version at iucnredlist.org. The authoritative source for the categories, criteria and assessment statistics used in Section 1.4.3.

CHAPTER II: Quantitative Wildlife Science and Monitoring

Chapter 1 established what wildlife populations are and how they behave. This chapter addresses a harder question: how do we know? Almost nothing a wildlife scientist wants to measure can be observed completely. Animals are cryptic, mobile, nocturnal, sparse or all four at once, so every field observation is a filtered and incomplete view of the underlying state. Quantitative wildlife science is the discipline of reasoning rigorously from that incomplete view to defensible statements about abundance, distribution, movement and risk. Section 2.1 develops the three estimation frameworks that dominate modern practice. Section 2.2 turns to the non-invasive technologies that now supply most survey data. Section 2.3 treats movement and telemetry, the basis of home range and connectivity analysis. Section 2.4 closes the loop by converting estimates into projections of extinction risk that can inform a decision.

2.1 Population Estimation Techniques

2.1.1 Counts, indices and estimates

Every field survey produces a count statistic C : the number of animals seen, heard, photographed or captured. That count is related to the abundance N of interest by a detection probability p :

$$E[C] = Np$$

This deceptively simple identity organises the whole of survey design (Figure 2.1). An index treats p as an unknown constant, then reports C or C per unit effort as a relative measure. An estimate models p explicitly, then reports N with an associated variance. Indices are cheaper and remain widespread, but they carry a specific and serious risk: any factor that changes p over time or space will masquerade as a change in abundance. Observer experience, wind speed, leaf cover, season, time of day, road access and even animal density itself all shift detection probability. Figure 2.1b shows a monitoring series in which abundance was constant while detection declined gently for fifteen years, producing an apparent decline of 38 per cent that would be indistinguishable from a real one.

An index is defensible in only two situations: when the analyst can argue that p is constant across every comparison being made; or when the index has been calibrated against an estimate obtained by another method. In practice this argument is rarely available, which is why the three frameworks that follow, all of which model detection, have become standard. Table 2.1 sets out their data requirements alongside the quantity each delivers.




Figure 2.1. The inference chain from field observation to population estimate. (a) Every count is the product of abundance and detection probability, so an analysis either assumes detection is constant (an index) or models it (an estimate). (b) A simulated fifteen-year monitoring series in which true abundance was constant while detection probability declined from 0.55 to 0.24. The raw count falls by 38 per cent, which a naive analysis would report as a population decline.

Table 2.1. The principal frameworks for estimating abundance and distribution, with their data requirements and outputs.

Framework	Field data required	Key assumptions	Primary output	Typical use
Distance sampling	Perpendicular or radial distance to each detection along random transects	Detection certain on the line; no responsive movement; distances measured accurately	Density and abundance with variance	Ungulates, seabirds, cetaceans, songbirds
Capture-recapture	Individual encounter histories over repeated occasions	Marks are permanent and read correctly; the model of capture heterogeneity is adequate	Abundance, survival, recruitment	Individually identifiable species
Spatial capture-recapture	Encounter histories with detector coordinates	Activity centres are fixed; detection declines with distance	Density, without an arbitrary sampled area	Camera-trapped and genetically sampled carnivores
Occupancy modelling	Repeat detection or non-detection at each site	Sites are closed within a season; detections are independent	Proportion of sites occupied; colonisation and extinction	Rare or cryptic species; large-scale monitoring
Index counts	Counts standardised by effort	Detection probability constant across comparisons	Relative abundance only	Long-running schemes; screening surveys

2.1.2 Distance sampling

Distance sampling exploits a simple insight: if animals are distributed independently of the transect line, then the decline in the number of detections with distance from the line reveals the detection process itself. The observer walks or flies a randomly placed line, recording the perpendicular distance x to every animal or cluster detected (Figure 2.2a). The distribution of those distances is used to fit a detection function $g(x)$, the probability of detecting an animal given that it is at distance x , scaled so that $g(0)$ equals one.

Two functional forms dominate. The half-normal function has a single scale parameter, while the hazard-rate function has an additional shape parameter that produces a broader shoulder near the line. Series expansions can be added to either. Model choice is normally made by the Akaike Information Criterion, subject to the shape criterion that a plausible detection function must have a shoulder: detection should not fall away immediately at the line (Figure 2.2c). A spiked distribution usually signals a field problem such as animals being flushed from the line rather than a genuine detection process.

The fitted function yields the effective strip half-width, the distance at which the number of animals missed inside equals the number detected outside:

$$\mu = \int g(x) dx, \text{ evaluated from } 0 \text{ to the truncation distance } w$$

Density then follows directly. For a line transect of total length L with n detections:

$$D = n / (2 L \mu)$$

and for point transects with k points and radial distances, the analogous estimator uses the effective detection radius. The practical requirements are three. First, the design must place lines at random with respect to the animals, normally by systematic random sampling with a random start, because non-random placement along tracks or ridgelines violates the assumption that animal distribution is independent of the line. Second, the largest 5 to 10 per cent of distances are usually truncated, since they contribute little information but strongly influence the fitted tail. Third, the survey needs enough detections: 60 to 80 for a coefficient of variation near 15 per cent, with more required where encounter rate varies strongly among transects (Figure 2.2e).




Figure 2.2. Distance sampling. (a) Randomly placed line transects with perpendicular distances to detections. (b) The observed distance distribution with half-normal and hazard-rate detection functions fitted and the effective strip half-width marked. (c) The shape criterion: a plausible detection function has a shoulder near the line. (d) If detection on the line is not certain, abundance is underestimated in direct proportion to the shortfall. (e) The relationship between the number of detections and the precision of the density estimate.

Table 2.2. Assumptions of conventional distance sampling, their consequences when violated and the standard remedies.

Assumption	Consequence of violation	Remedy
Detection on the line is certain, $g(0) = 1$	Abundance underestimated in proportion to the shortfall (Figure 2.2d)	Double-observer or mark-recapture distance sampling to estimate $g(0)$
Animals do not move in response to the observer	Attraction inflates density; avoidance deflates it	Detect animals before they respond; use aerial or acoustic platforms
Distances are measured without error	Rounding to convenient values distorts the fitted function	Use rangefinders; group into bins; check for heaping in the histogram
Detections are independent	Variance underestimated where animals occur in clusters	Record cluster size; estimate expected cluster size by regression on distance
Transects are placed at random relative to animals	Systematic bias of unknown sign	Systematic random design with a random start; never survey along existing tracks

2.1.3 Capture-recapture

Where individuals can be marked and re-encountered, capture-recapture uses the overlap between samples to estimate the fraction of the population that was seen. In its simplest closed-population form, with n_1 individuals marked on the first occasion, n_2 caught on the second of which m_2 carry marks, the Lincoln-Petersen estimator follows from the proportionality argument that the marked fraction of the second sample estimates the marked fraction of the population:

$$\hat{N} = n1 \ n2 / m2$$

This estimator is biased upward at small sample sizes, so the Chapman modification is preferred in practice:

$$\hat{N} = (n1 + 1)(n2 + 1) / (m2 + 1) - 1$$

With more than two occasions, the encounter history of each individual becomes a string of ones and zeros (Figure 2.3a) that supports a much richer family of models. Closed-population models, which assume no births, deaths, immigration or emigration during sampling, are distinguished by the sources of variation in capture probability they permit (Table 2.3). Failing to model heterogeneity is the most common error, because individuals that are intrinsically hard to catch are systematically underrepresented, which biases abundance downward.

Table 2.3. The standard closed-population capture-recapture model set. Models may be combined, giving Mth, Mbh and so on.

Model	Source of variation in capture probability	Biological interpretation
M0	None; capture probability constant	Rarely realistic; a null model for comparison
Mt	Time; probability varies among occasions	Weather, effort or observer differs between occasions
Mb	Behaviour; probability changes after first capture	Trap-happy animals seeking bait; trap-shy animals avoiding devices
Mh	Individual heterogeneity	Differences in home range position, sex, age or boldness
Mth, Mbh, Mtbh	Combinations of the above	Realistic for most field data; requires larger samples

Open-population models relax the closure assumption and estimate demographic rates directly. The Cormack-Jolly-Seber model estimates apparent survival together with recapture probability, where apparent survival confounds mortality with permanent emigration. The Jolly-Seber family adds abundance and recruitment. Pollock's robust design combines the two by nesting several closely spaced secondary occasions, treated as closed, within widely spaced primary periods, treated as open, which allows temporary emigration to be separated from mortality.

The persistent weakness of conventional capture-recapture is that it estimates abundance rather than density, because the area effectively sampled by a trap array is unknown. Animals whose home ranges lie partly outside the array can still be captured, so the effective sampled area exceeds the geometric area of the array. Dividing abundance by the array area therefore overestimates density. The size of the error grows with the movement scale of the species relative to the array (Figure 2.3c). Spatial capture-recapture resolves the problem by modelling the location of each individual's activity centre as a latent variable, with detection declining as a function of distance from that centre (Figure 2.3b):

$$p(d) = g_0 \exp(-d^2 / 2\sigma^2)$$

Here g_0 is the detection probability at the activity centre while σ is a movement scale parameter closely related to home range radius. Because the model estimates the density of activity centres over a defined state space, it returns density directly, with no arbitrary buffer. Spatial capture-recapture is now the default for camera-trap and genetic surveys of individually recognisable carnivores.




Figure 2.3. Capture-recapture. (a) An encounter history matrix, the raw data for all capture-recapture models. (b) Spatial capture-recapture: each individual has a latent activity centre, with detection probability declining with distance from it, so detector locations become part of the model. (c) The edge effect: dividing abundance by the geometric area of the detector array overestimates density, increasingly so as the movement scale of the species grows relative to the array.

2.1.4 Occupancy modelling

For species too sparse or too cryptic to yield adequate capture samples, the practical state variable is not abundance but occupancy: the proportion of sites at which the species is present. The central difficulty is that a non-detection is ambiguous, since it may mean the species is absent or merely that it was missed. Repeat visits resolve the ambiguity. The single-season occupancy model has two parameters: ψ , the probability that a site is occupied; and p , the probability of detecting the species at an occupied site during one visit. For a site surveyed K times with detection history h , the likelihood contribution is

$$\psi \prod p(k)^{h(k)} (1 - p(k))^{(1-h(k))} \text{ for sites with a detection}$$

and for a site with no detection at all, the contribution is ψ times the probability of missing the species K times added to the probability that the site was unoccupied:

$$\psi (1 - p)^K + (1 - \psi)$$

Both ψ and p can be modelled as functions of covariates through a logit link, which is what makes the framework so useful: habitat covariates enter the occupancy component while survey covariates such as observer, weather or effort enter the detection component. Naive occupancy, the raw proportion of sites with a detection, is always biased low by exactly the factor 1 minus

the probability of detecting the species at least once (Figure 2.4c). With p of 0.15 and four visits, naive occupancy recovers barely half of true occupancy.




Figure 2.4. Occupancy modelling. (a) A detection history matrix. Sites that are occupied but never detected are the reason repeat visits are needed. (b) The probability of detecting a species at least once rises with the number of surveys, at a rate governed by per-visit detection probability. (c) Naive occupancy is biased low; the shaded region is the bias for a species with per-visit detection of 0.15.

Survey design therefore turns on a trade-off between the number of sites and the number of visits per site. For a fixed budget, rare species with low detection probability require more visits per site, whereas widespread and easily detected species are better served by more sites with fewer visits (Table 2.4). The general rule is that at least three visits are needed to estimate p at all. The standard error of ψ is minimised near the effort levels shown in the table.

Table 2.4. Indicative number of survey visits per site required to detect a species at an occupied site with probability 0.95, given per-visit detection probability. Values are the smallest K satisfying $1 - (1 - p)^K \geq 0.95$.

Per-visit detection p	K for 0.80	K for 0.95	K for 0.99	Typical example
0.10	16	29	44	Cryptic snake on a timed search
0.20	8	14	21	Small mammal in a single trap night
0.30	5	9	13	Amphibian call survey, poor conditions
0.50	3	5	7	Camera trap over a one-week deployment
0.70	2	3	4	Territorial songbird in the breeding season

Two extensions matter for management. Multi-season occupancy models, also called dynamic models, add colonisation and local extinction probabilities between seasons, which converts the

framework into a direct test of metapopulation dynamics as described in Section 1.2.5. The Royle-Nichols model exploits the fact that detection probability at a site normally rises with local abundance, so it can extract an index of abundance from detection or non-detection data alone.

2.2 Non-Invasive Survey Technologies

The technologies in this section share a common advantage: they gather data without capturing, handling or in most cases even disturbing the animal. That removes a welfare cost, removes the capture bias associated with trapping and allows sampling at scales that handling could never reach. Their common limitation is that the raw output is a detection rather than an animal, so each requires an explicit analytical bridge to the population quantity of interest. Table 2.5 compares them.

Table 2.5. Non-invasive survey technologies compared. Cost bands are indicative of equipment plus analysis effort per unit of area surveyed.

Technology	Best suited to	Direct output	Route to abundance	Principal limitation
Camera trapping	Medium and large terrestrial mammals; ground birds	Time-stamped images at fixed points	SECR if individuals are recognisable; REM or camera-trap distance sampling if not	Detection zone varies with target size, temperature and vegetation
Passive acoustic monitoring	Vocal birds, bats, anurans, cetaceans, insects	Continuous audio and derived detections	Cue-rate models; occupancy; acoustic SECR with arrays	Requires species-specific recognisers and validation
Environmental DNA	Aquatic and semi-aquatic taxa; whole community surveys	Presence of target sequences in a sample	Occupancy with hierarchical replication; abundance only weakly	Transport and degradation blur the location and timing of presence
Non-invasive genetics	Carnivores and ungulates via scat, hair or shed feathers	Individual genotypes and sex	SECR or closed capture-recapture on genotypes	Genotyping error creates false individuals
Thermal and drone imagery	Open-habitat aggregations; colonial nesters	Georeferenced counts	Distance sampling or total counts with detection correction	Weather dependent; regulatory restrictions on flight

2.2.1 Camera trapping

A camera trap converts a point in the landscape into a continuously operating observer. Design begins with the same principle as any survey: cameras must sample the study area in a way that supports the intended inference. Placing cameras on trails and at scent stations maximises the number of images but produces a biased sample, because trail use differs among species, sexes and individuals. Randomised or systematic placement, with camera spacing set relative to the home range diameter of the focal species, is required whenever a density estimate is the goal.

For species with individually distinctive markings, spatial capture-recapture applied to camera data is the method of choice. For unmarked species the two viable routes are camera-trap distance sampling, which uses the distance of each animal from the camera at a fixed set of snapshot moments; and the random encounter model, which relates the trapping rate to density through the geometry of the detection zone and the average movement speed of the animal (Figure 2.5a):

$$D = (y/t) \times \pi / (v r (2 + \theta))$$

where y is the number of independent encounters in time t , v is the average day range of the species, r is the radius of the detection zone and θ is its angle in radians. The model requires an independent estimate of v , which is normally obtained from telemetry, so the random encounter model is not free of animal handling despite the technology being non-invasive.

A common but weak practice is to report a relative abundance index as events per 100 camera-nights. Its central problem is that the analyst must choose an independence interval to decide when consecutive images represent a new event. That choice is arbitrary. The resulting index varies severalfold across the range of intervals in routine use (Figure 2.5c). Reporting such an index without stating the interval and without evidence that detection probability is comparable across the units being compared, is not defensible.




Figure 2.5. Camera trapping. (a) The geometry of the random encounter model: density is recovered from the trapping rate given the detection radius r , the detection angle θ and the average day range v . (b) Species accumulation with survey effort, showing why detecting the full community requires far more effort than detecting the common species. (c) The relative abundance index computed from one set of images under eight different independence intervals.

2.2.2 Bioacoustic monitoring

Autonomous recording units allow continuous acoustic sampling at many sites for weeks or months, which suits vocal taxa that are otherwise hard to survey. Two configuration decisions dominate data quality. The sampling rate must be at least twice the highest frequency of interest, which means 44 kHz or higher for most birds and 250 kHz or higher for echolocating bats. The duty cycle then trades battery life and storage against temporal coverage. It must be set with the diel activity pattern of the target species in mind (Figure 2.6c), since a schedule that samples only at dawn will systematically miss a crepuscular or nocturnal caller.

Analysis converts audio into a time-frequency representation, the spectrogram (Figure 2.6a), on which species-specific signals are identified either by a template-matching recogniser or by a convolutional neural network trained on labelled examples. Every recogniser embodies a threshold that trades precision against recall (Figure 2.6b). Precision is the proportion of flagged detections that are genuine; recall is the proportion of genuine calls that are flagged. A high threshold yields few false positives at the cost of missing many calls, while a low threshold does the reverse. Neither error is inherently worse: the appropriate operating point depends on the analysis. False positives are especially damaging in occupancy analysis, because a single false detection at an unoccupied site inflates the occupancy estimate, so most workflows either verify a sample of detections manually or use occupancy models that explicitly accommodate false positives.

Acoustic data support several levels of inference. Detection or non-detection feeds directly into occupancy models. Where the cue rate, meaning calls per individual per unit time, can be estimated independently, call counts can be converted into density. Arrays of synchronised recorders allow acoustic localisation of the caller, which supports spatial capture-recapture. Community-level acoustic indices that summarise the whole soundscape are attractive because they require no species identification, but their relationship to biodiversity is inconsistent across systems, so they should be treated as exploratory rather than as a substitute for species-level analysis.




Figure 2.6. Bioacoustic monitoring. (a) A spectrogram showing four frequency-modulated calls with their harmonics. (b) The precision-recall trade-off inherent in any automated recogniser:

moving the score threshold trades false positives against missed calls. (c) A diel calling pattern with crepuscular peaks, which determines an appropriate recording schedule.

2.2.3 Environmental DNA

All organisms shed DNA into their surroundings in mucus, faeces, gametes, urine and sloughed cells. Environmental DNA methods recover that material from a water, soil or air sample, then identify the species that contributed it. Two analytical modes are distinguished. Targeted assays use species-specific primers with quantitative or droplet digital PCR to test for one species at high sensitivity. Metabarcoding uses universal primers to amplify a taxonomically informative region across a whole assemblage, then assigns the resulting sequences by matching them against a reference database.

The workflow (Figure 2.7a) has more failure modes than most field methods. Every one of them must be controlled. False negatives arise from low target concentration, from PCR inhibitors such as humic acids, from primer mismatch and from DNA degradation, which proceeds within hours to weeks depending on temperature, ultraviolet exposure and microbial activity. False positives arise from field and laboratory contamination, from tag-jumping between samples on a sequencing run and from incomplete or erroneous reference databases. Table 2.6 sets out the standard controls.

Table 2.6. Sources of error in environmental DNA surveys and the controls that address them.

Error	Mechanism	Detection	Control
False negative	Low target concentration; inhibition; primer mismatch	Internal positive control fails to amplify	Larger sample volume; inhibition removal; replicate PCRs; validated primers
False positive, field	Carry-over on nets, boots or bottles	Field blank amplifies	Single-use or bleached equipment; a field blank at every site
False positive, laboratory	Aerosol or pipette contamination	Extraction and PCR blanks amplify	Separate pre-PCR and post-PCR rooms; unidirectional workflow
Misassignment	Incomplete or erroneous reference database	Implausible species in the output	Curated regional reference library; conservative identity thresholds
Spatial or temporal ambiguity	Transport downstream; persistence after the animal has left	Detection without any other evidence of presence	Sample near suspected sources; interpret as presence within a catchment, not at a point

Because a site yields several samples while each sample yields several PCR replicates, environmental DNA data have a natural hierarchical structure that maps onto a multi-scale occupancy model. Occupancy at the site is ψ , the probability that a sample from an occupied site contains target DNA is θ ; and the probability that a PCR replicate from a positive sample

amplifies is p . Estimating all three separates true absence from failure to capture DNA from failure to amplify it, which is the single most important methodological advance in the field (Figure 2.7b). Reporting a raw list of detections without this structure discards the information needed to interpret it.

2.2.4 Non-invasive genetic sampling

Faeces, hair, feathers, saliva on kill sites and shed skin all carry nuclear DNA that can be genotyped to identify individuals, sex and relatedness. Applied to hair snares or systematic scat transects, this converts an unmarked population into a marked one, so the full capture-recapture and spatial capture-recapture toolkit becomes available. The approach has transformed the monitoring of bears, wolves, wolverines and many primates.

The technical obstacle is that non-invasively collected DNA is degraded and present in low quantity, which causes two characteristic errors. Allelic dropout occurs when one allele of a heterozygote fails to amplify, so a heterozygote is scored as a homozygote. False alleles arise from PCR artefacts. Both create genotypes that differ slightly from the true one, so a single individual can be recorded as several, which inflates the apparent population size. The standard defence is the multiple-tubes approach: each sample is genotyped several times independently, with a consensus required before a genotype is accepted (Figure 2.7c). Quality thresholds should be reported alongside the estimate, as should the probability of identity, meaning the probability that two distinct individuals share a genotype by chance.

Faecal samples carry more than genotypes. Diet can be reconstructed by metabarcoding the prey DNA they contain, parasites can be identified directly, while steroid metabolites provide non-invasive measures of reproductive state and of physiological stress. A single well-designed scat survey can therefore deliver density, diet, disease and stress data from one collection effort, which is why it has become a mainstay of carnivore monitoring.




Figure 2.7. Molecular non-invasive methods. (a) The environmental DNA workflow, with blanks required at every stage. (b) Detection probability at an occupied site as a function of the number of PCR replicates, for three values of the sample-level detection probability θ . (c) The probability of

recovering a correct eight-locus genotype from degraded faecal DNA, as a function of the number of independent replicates and the per-allele dropout rate.

2.3 Movement Ecology and Telemetry

Movement links the individual to the landscape. It determines which resources an animal can reach, which mates it can find, which risks it encounters and which populations it connects. The movement ecology framework organises any trajectory around four components: the internal state that motivates movement, the motion capacity that permits it, the navigation capacity that orients it and the external factors that constrain it. Telemetry provides the empirical record.

2.3.1 Telemetry technologies

Very high frequency radio telemetry remains in use because tags are light, cheap and long-lived, but each location requires an observer with a receiver, so sample sizes are small and biased toward accessible terrain. Argos satellite transmitters give global coverage with location errors of hundreds of metres to several kilometres. GPS tags give errors of a few tens of metres or less, either storing locations for later recovery or transmitting them over cellular or satellite networks. Light-level geolocators are extremely light and therefore suited to small migratory birds, but their errors are on the order of one to two hundred kilometres, which restricts them to continental-scale questions. Accelerometers and other biologging sensors add behavioural and physiological detail that location data alone cannot provide.

Every deployment involves the same three-way trade-off among tag mass, fix frequency and duration (Figure 2.8). Battery capacity scales with mass, while energy consumption scales with the number of fixes, so a fine-grained record is necessarily a short one unless the tag is heavy. The mass constraint is not negotiable: the widely used guideline caps tag mass at 5 per cent of body mass for terrestrial mammals and 3 per cent for birds, with lower limits for aerial foragers and for species that must escape predators by flight. Table 2.7 summarises the main options.

Table 2.7. Telemetry technologies compared. Values are indicative and change rapidly as hardware improves.

Technology	Location error	Minimum mass	Typical duration	Principal use
VHF radio	50 to 500 m	Under 1 g	Weeks to a year	Small species; survival monitoring; homing to a den or nest
Argos satellite	250 m to 5 km	5 to 20 g	Months to a year	Marine and long-distance migrants beyond cellular coverage
GPS store-on-board	5 to 20 m	10 to 30 g	Months	Fine-scale behaviour where the tag can be recovered
GPS with cellular or satellite link	5 to 20 m	20 to 60 g	One to five years	Large mammals; near real-time conflict and poaching alerts
Light-level geolocator	50 to 200 km	0.3 to 1.5 g	One annual cycle	Migration routes of small birds; requires recapture
Accelerometer and biollogger	Not applicable	1 to 20 g	Days to months	Behavioural state, energy expenditure, foraging events

Two data quality problems recur. Fix success is not random: GPS tags fail to acquire a fix more often under dense canopy, in steep terrain and when the animal is resting in a burrow or den, which biases the recorded distribution toward open habitat. The remedy is to model fix success as a function of habitat, then weight locations accordingly. Location error is likewise structured rather than random, so analyses that treat locations as exact will overstate the precision of any derived quantity. State-space models that separate the true track from observation error are the standard solution.




Figure 2.8. Telemetry trade-offs. (a) Location error against deployment duration for the main technologies, with symbol area proportional to typical tag mass. (b) The duty-cycle trade-off: shortening the fix interval increases behavioural resolution while shortening the record, because battery capacity is fixed by tag mass.

2.3.2 Home range estimation

A home range is the area an animal traverses in its normal activities of foraging, mating and caring for young. Making that definition operational has produced a long succession of estimators, each embodying different assumptions (Table 2.8).

The minimum convex polygon is the smallest convex polygon containing all locations. It is simple, comparable across decades of literature and defensible as a crude measure of extent, but it is highly sensitive to outliers, it includes large areas the animal never uses, while its area increases without limit as more locations are added. Removing the most peripheral 5 per cent of points gives the 95 per cent minimum convex polygon, which is more stable but arbitrary.

Kernel density estimation instead builds a utilisation distribution, a probability density describing the relative time spent at each point in space. Contours of that surface then delimit the home range, conventionally at 95 per cent and with the 50 per cent contour taken as the core area. The critical choice is the smoothing bandwidth h . The reference bandwidth assumes an underlying bivariate normal distribution and typically oversmooths, merging distinct activity centres. Least squares cross-validation often undersmooths, fragmenting the range into isolated islands. Figures 2.9b, 2.9c and 2.9e show how strongly the estimate depends on this single parameter, so the bandwidth and the method used to select it must always be reported.

Both estimators assume that locations are independent, an assumption that modern GPS data violate flagrantly: successive fixes minutes apart are strongly autocorrelated. Autocorrelated kernel density estimation addresses this directly by first fitting a continuous-time movement model to the data, then using the fitted model to weight the locations appropriately. It typically returns a larger area than conventional kernel density estimation together with an honest confidence interval, which conventional estimators do not provide (Figure 2.9f). A related development, the Brownian bridge movement model, uses the time ordering of locations to estimate the probability of occupancy between fixes, which makes it well suited to migration corridors where the animal is in transit rather than resident.

Whatever the estimator, the area-observation curve should be inspected (Figure 2.9d). Plotting the estimated area against the number of locations shows whether sampling has been sufficient for the estimate to stabilise. A curve still rising steeply at the final sample size indicates that the animal has not been followed long enough for the home range to be defined.




Figure 2.9. Home range estimation applied to one simulated track of 900 locations. (a) Minimum convex polygons at 100 and 95 per cent. (b) and (c) Kernel density estimates under a small and a reference bandwidth, with 50 and 95 per cent contours. (d) The area-observation curve, used to judge whether sampling is adequate. (e) Sensitivity of the estimated area to the bandwidth. (f) The five estimators compared on the same data, with the interval on the autocorrelated estimate showing the uncertainty the others conceal.

Table 2.8. Home range estimators, their assumptions and their appropriate use.

Estimator	What it estimates	Key assumption	When to use
Minimum convex polygon	Outer extent of recorded locations	None, but assumes convexity is meaningful	Comparison with historical studies; crude extent only
Kernel density estimation	Utilisation distribution	Locations are independent	Well-spaced VHF locations; report the bandwidth method
Autocorrelated KDE	Utilisation distribution with autocorrelation modelled	A continuous-time movement model fits the data	GPS data with short fix intervals; the current default
Brownian bridge movement model	Probability of occupancy between fixes	Movement between fixes is a random walk	Migration routes; corridor identification
Local convex hull	Utilisation distribution respecting hard boundaries	Sharp edges are real	Riparian, coastal or fragmented habitat

2.3.3 Behavioural states and corridors

A trajectory can be decomposed into step lengths and turning angles. The joint distribution of those two quantities distinguishes behavioural modes: short steps with uniformly distributed turns indicate area-restricted search or resting, while long steps with turns concentrated near zero indicate directed travel. Hidden Markov models formalise this by treating the behavioural state as an unobserved variable that switches according to a transition matrix, with state-dependent distributions for step length and turning angle (Figure 2.10). The fitted model returns the most likely state sequence, which converts a raw track into a behavioural time series that can then be related to habitat, season or human disturbance.




Figure 2.10. Decomposing a trajectory. (a) A simulated track decoded by a two-state hidden Markov model into encamped and travelling segments. (b) State-dependent step length distributions. (c) Turning angle distributions: near-uniform when encamped, concentrated near zero when travelling.

Connectivity analysis asks a different question: which parts of the landscape carry movement between habitat patches? The standard workflow begins with a resistance surface, a raster in which each cell is assigned a cost of movement. Resistance values should be derived empirically wherever possible, most often from a step selection function fitted to telemetry data, which compares the habitat at each used step endpoint against the habitat at a set of available alternatives (Figure 2.11c). Expert opinion remains common but produces surfaces whose sensitivity to the assumed values is rarely tested.

Two families of algorithm then operate on the surface (Table 2.9). Least-cost path analysis finds the single route of minimum accumulated cost between two patches, which is appropriate when the animal has good knowledge of the landscape. Circuit theory treats the landscape as a conductive surface across which current flows, so it identifies all routes in proportion to their conductance, which is appropriate when movement is exploratory rather than informed. The two approaches make different predictions in the same landscape. The difference matters: circuit theory identifies pinch points where a large flux is squeezed through a narrow passage (Figure 2.11b), which are precisely the locations where a crossing structure or an easement delivers most benefit per unit cost. Whichever method is used, the output is a hypothesis that should be validated against independent movement or genetic data before it is used to commit funds.




Figure 2.11. Corridor identification. (a) A resistance surface with a road bisecting the landscape and an underpass at its centre, with the least-cost path from patch A to patch B. (b) Circuit-theory current density on the same surface, identifying the underpass as a pinch point through which a large movement flux is concentrated. (c) Coefficients from a step selection function, the preferred empirical basis for assigning resistance values, with 95 per cent confidence intervals.

Table 2.9. Approaches to modelling landscape connectivity.

Approach	Behavioural assumption	Output	Best use
Euclidean distance	The matrix is uniform	Straight-line distance between patches	A null model for comparison
Least-cost path	The animal has complete knowledge and takes the optimal route	A single line of least accumulated cost	Species with large home ranges and good spatial memory
Circuit theory	Movement is a random walk weighted by conductance	Current density across all possible routes	Dispersing juveniles; identifying pinch points
Individual-based simulation	Explicit movement rules with mortality and perception	Simulated trajectories and success rates	Where behaviour is well known and stakes are high
Landscape genetics	Gene flow integrates realised movement over generations	Correlation of genetic with landscape distance	Validating any of the above against realised connectivity

2.4 Data Analysis and Population Viability

2.4.1 Detecting trends

The most common question asked of a monitoring programme is whether the population is declining. Answering it requires more than a series of estimates, because any series contains two distinct kinds of variation. Observation error is the difference between the estimate and the true state at that moment; process variation is real change in the true state driven by weather, resources and demographic chance. Confusing the two leads directly to error: treating observation error as process variation exaggerates apparent volatility and inflates estimated extinction risk, while ignoring observation error entirely produces overconfident trend estimates. State-space models separate them by specifying a process equation for the latent state alongside an observation equation linking the state to the data (Figure 2.12b).

Before a monitoring programme begins, its statistical power should be evaluated: the probability that it will detect a decline of a specified size if one is occurring. Power rises with the number of years, with the magnitude of the decline and with the precision of each estimate. It falls sharply as the coefficient of variation of individual estimates rises. Figure 2.12a shows the practical consequence. A programme whose annual estimates carry a coefficient of variation of 0.40 needs roughly twice as long to reach 80 per cent power as one achieving 0.10. Many long-running schemes have been unable to detect biologically important declines for this reason alone, which is why power analysis belongs at the design stage rather than in the discussion of the final report.




Figure 2.12. Monitoring design and analysis. (a) Statistical power to detect a 5 per cent annual decline as a function of monitoring duration, for three levels of precision in the annual estimates. (b) A state-space model separates the latent population trajectory from observation error; the vertical lines show the residual observation error at each survey.

2.4.2 Population viability analysis

Population viability analysis projects a population forward under explicit assumptions about vital rates and their variability, then summarises the resulting distribution of outcomes as a probability of falling below a specified threshold within a specified period. Three model classes are in general use (Table 2.10). Count-based analysis uses only a time series of abundance, from which the mean and variance of the log growth rate are estimated and used in a diffusion approximation. It is attractive because the data requirement is minimal, but it cannot represent age structure, density dependence or management actions that target one vital rate. Demographic analysis uses a matrix model, so it can evaluate stage-specific interventions and supports the elasticity analysis introduced in Section 1.2.4. Individual-based analysis simulates each animal separately, which permits explicit treatment of mating systems, inbreeding, dispersal among patches and behavioural rules, at the cost of many more parameters.

Table 2.10. Approaches to population viability analysis.

Approach	Data required	What it can represent	Main limitation
Count-based, diffusion approximation	A time series of abundance, ideally 10 years or more	Mean growth rate with its variance	No age structure; cannot evaluate stage-specific management
Demographic, matrix-based	Age or stage-specific survival and fecundity with their variances	Stage structure, elasticity, targeted interventions	Vital rate variances are hard to estimate and often guessed
Individual-based simulation	Full demography plus behavioural and genetic parameters	Mating system, inbreeding, dispersal, spatial structure	Parameter-hungry; apparent realism can mask deep uncertainty
Metapopulation PVA	Patch-specific demography plus dispersal rates	Regional dynamics, rescue effects, correlated catastrophes	Dispersal rates are the least well known parameters

Whatever the class, the analysis must represent the four sources of variation set out in Table 2.11. Demographic stochasticity is the chance variation in survival and reproduction among individuals. It matters only in small populations, because the law of large numbers averages it away as numbers rise, but it is precisely the small populations that concern conservation. Environmental stochasticity is year-to-year variation in vital rates driven by external conditions, affecting populations of all sizes. Catastrophes are rare events of large effect, such as an epidemic, a hurricane or a wildfire; because they are rare, they are poorly characterised, yet they frequently dominate the risk estimate. Genetic stochasticity, as described in Section 1.4.5, erodes adaptive potential and depresses vital rates through inbreeding.

Table 2.11. Sources of stochasticity in population projections and their relative importance.

Source	Mechanism	Depends on population size?	How it enters a model
Demographic	Chance variation among individuals in survival and reproduction	Strongly; negligible above a few hundred	Binomial and Poisson draws per individual
Environmental	Year-to-year variation in conditions affecting all individuals	No	Random draws of vital rates from a distribution across years
Catastrophic	Rare events of large effect: disease, fire, storm, spill	No	A probability per year with a multiplier on survival
Genetic	Drift and inbreeding depression in small populations	Strongly	Vital rates reduced as a function of the inbreeding coefficient
Parameter uncertainty	Imperfect knowledge of the true vital rates	No	Sampling from the posterior or by sensitivity analysis

Two outputs deserve emphasis. The quasi-extinction threshold is the abundance below which a population is regarded as functionally lost, commonly set at 50 or at the number below which the model is no longer credible. Using a threshold rather than zero avoids projecting into a range where the model assumptions certainly fail. The cumulative quasi-extinction curve then gives the probability of falling below that threshold by each future year. It is the natural currency for comparing management options (Figure 2.13b). Sensitivity analysis, in which each parameter is varied in turn while the others are held constant, identifies the parameters that most influence the result (Figure 2.13c). Its value is diagnostic as much as predictive: a parameter with high influence and high uncertainty is the first target for further fieldwork.




Figure 2.13. Population viability analysis of a hypothetical population of 260 individuals with a mean log growth rate of -0.012 , environmental standard deviation of 0.145 and a 2 per cent annual catastrophe probability. (a) Three hundred stochastic projections with the median and the 5th and

95th percentiles. (b) Cumulative probability of quasi-extinction under three scenarios. (c) Sensitivity of the 50-year risk to a 10 per cent change in each parameter.

2.4.3 Minimum viable population

The minimum viable population is the smallest population that has a specified probability of persisting for a specified time, conventionally 95 per cent over 100 years. The definition makes the concept explicitly conditional: an MVP figure is meaningless without the probability, the time horizon and the assumed conditions attached to it.


Two lines of evidence converge on similar numbers. The genetic argument, introduced in Section 1.4.5, holds that an effective size of at least 100 is needed to limit inbreeding depression in the short term while at least 1000 is needed to retain evolutionary potential. Since effective size is typically about one tenth of census size, this implies census populations in the thousands. The demographic argument, based on published viability analyses across many taxa, arrives at median MVP estimates in the low thousands of adults, with wide variation among species and considerable dependence on the modelling assumptions.

Three cautions apply. First, MVP is not a target: a population at its minimum viable size still carries a substantial probability of loss, so managing toward that number is managing toward an unacceptable risk. Second, generic MVP figures should not substitute for species-specific analysis, because the variation among taxa is larger than the central tendency. Third, an MVP figure derived under current conditions does not hold under changed ones, so any figure quoted in a recovery plan should be revisited as the environment changes. The defensible use of the concept is as a rough lower bound that flags when a population is unambiguously too small, not as a specification of how small a population may safely be allowed to become.

2.4.4 From estimates to decisions

Monitoring that does not change any decision is expenditure without benefit. Structured decision making addresses this by requiring that objectives, alternatives and the measures that discriminate among them are specified before data are collected. The value of information can then be calculated formally: how much would the expected outcome improve if a given uncertainty were resolved? Uncertainties that cannot change the chosen action have no value of information, however scientifically interesting they are, so this calculation is an efficient filter on monitoring priorities.

Adaptive management embeds that logic in a repeating cycle (Figure 2.14). Objectives and decision triggers are agreed first, so that the estimate a survey will produce is known in advance to be capable of changing the action. The survey is then designed to deliver the precision that the trigger requires, not merely the precision that the budget allows. Data are collected with detection modelled explicitly, state and trend are estimated, risk is projected, an action is taken, then the model is revised against the outcome. Each pass reduces uncertainty, so the sequence of decisions improves even where the initial understanding was poor. This is the framework within which every technique in this chapter finds its purpose. The chapters that follow assume it throughout.



Each pass reduces uncertainty, so later decisions rest on better estimates.

Figure 2.14. The adaptive management cycle. Monitoring is designed backwards from the decision it must inform, so the required precision is set by the decision trigger rather than by convenience. Each completed pass revises the model, which reduces the uncertainty facing the next decision.

Chapter Summary

- Every count is the product of abundance and detection probability. An index assumes detection is constant, which is rarely defensible, so modern practice models detection explicitly.
- Distance sampling recovers density from the decline in detections with distance from a randomly placed transect. Its critical assumptions are certain detection on the line, no responsive movement and accurate distance measurement.
- Capture-recapture estimates abundance from the overlap among samples. Unmodelled individual heterogeneity biases estimates downward, while spatial capture-recapture removes the arbitrary sampled area that troubles conventional models.
- Occupancy modelling uses repeat visits to separate absence from non-detection. Naive occupancy is always biased low, with the required number of visits rising steeply as per-visit detection falls.
- Camera trapping supports rigorous density estimation: through spatial capture-recapture for marked species; through the random encounter model or camera-trap distance sampling for unmarked ones. Relative abundance indices depend on an arbitrary independence interval.

- Bioacoustic recognisers trade precision against recall. False positives are particularly damaging in occupancy analysis, so a validation protocol is part of the method rather than an optional extra.
- Environmental DNA has a hierarchical error structure across sites, samples and PCR replicates, so it should be analysed with a model that reflects it. Blanks at every stage are mandatory.
- Telemetry involves an unavoidable trade-off among tag mass, fix frequency and duration. Fix success and location error are both structured rather than random, so both should be modelled.
- Minimum convex polygons measure extent while kernel methods estimate a utilisation distribution. Bandwidth choice dominates the result, so autocorrelated estimators are now preferred for GPS data.
- Connectivity models convert a resistance surface into predicted routes. Least-cost paths assume informed movement while circuit theory assumes exploratory movement, so only the latter identifies pinch points.
- Population viability analysis compares management options under stated assumptions rather than predicting the future. Quasi-extinction curves and sensitivity analyses are its most useful outputs.
- Minimum viable population figures are conditional on a probability, a time horizon and a set of conditions. They flag populations that are unambiguously too small; they are not management targets.

Key Terms

Autocorrelated KDE a home range estimator that fits a continuous-time movement model before weighting locations

Detection function $g(x)$ the probability of detecting an animal at distance x from the transect, scaled so that $g(0) = 1$

Detection probability p the probability that an animal or species present is recorded by a survey

Effective strip half-width the distance at which animals missed inside equal animals detected outside

Encounter history the sequence of detections and non-detections recorded for one individual or site

Environmental DNA genetic material recovered from water, soil or air rather than from the organism

Hidden Markov model a model that infers unobserved behavioural states from step lengths and turning angles

Naive occupancy the raw proportion of sites with at least one detection; always biased low

Occupancy ψ the probability that a site is occupied by the focal species

Pinch point a location where circuit theory predicts a large movement flux through a narrow passage

Precision and recall the proportion of flagged detections that are genuine; the proportion of genuine events flagged

Quasi-extinction threshold the abundance below which a population is treated as functionally lost

Random encounter model a method converting camera trapping rate to density using detection geometry and day range

Resistance surface a raster assigning a cost of movement to each cell of the landscape

Spatial capture-recapture a model estimating density from encounter histories with detector coordinates

Utilisation distribution a probability density describing the relative time an animal spends at each location

Review Questions

1. A line transect survey covers 40 km and records 96 detections. The fitted half-normal detection function gives an effective strip half-width of 62 m. Estimate density in animals per square kilometre, then state which assumption you would test first.
2. A camera-trap study of an unmarked ungulate reports 18 events per 100 camera-nights in a protected area and 7 in an adjacent hunting concession. List three explanations besides a genuine difference in abundance, then describe the analysis you would run instead.
3. Explain why the naive proportion of occupied sites is always biased low. For a species with per-visit detection of 0.25, calculate the number of visits needed to be 95 per cent confident of detecting it at an occupied site.
4. A conventional capture-recapture study divides estimated abundance by the geometric area of the trap array. Explain the direction of the resulting bias, then explain how spatial capture-recapture removes it.
5. An environmental DNA survey detects a target fish at 3 of 20 sites, with 2 of 8 PCR replicates positive at each. Explain how a multi-scale occupancy model would use these data, then identify the controls you would require before accepting the result.
6. A colleague reports a home range of 412 hectares from a kernel density estimate. What three pieces of information do you need before you can interpret that number?
7. Contrast least-cost path analysis with circuit theory. For a juvenile dispersing through unfamiliar terrain, which is the more appropriate model? Explain why.
8. A population viability analysis gives a 22 per cent probability of quasi-extinction within 50 years, with the result most sensitive to adult survival. Outline the monitoring and management response that follows.

Further Reading

- Buckland, S. T. et al. Distance Sampling: Methods and Applications. The standard reference for the material in Section 2.1.2.
- Williams, B. K., Nichols, J. D. and Conroy, M. J. Analysis and Management of Animal Populations. Comprehensive treatment of estimation and decision analysis.
- MacKenzie, D. I. et al. Occupancy Estimation and Modeling. Full development of the framework in Section 2.1.4, including design guidance.
- Royle, J. A. et al. Spatial Capture-Recapture. The definitive account of the models introduced in Section 2.1.3.
- Fleming, C. H. and Calabrese, J. M. Papers on autocorrelated kernel density estimation and the ctmm package, for the methods in Section 2.3.2.
- Morris, W. F. and Doak, D. F. Quantitative Conservation Biology. Accessible development of count-based and demographic viability analysis.

CHAPTER III: HABITAT ECOLOGY, MANAGEMENT AND RESTORATION

Habitat is where the theory of Chapter 1 and the measurement of Chapter 2 meet the ground. Almost every management action available to a wildlife agency operates on habitat rather than on animals directly: a burn is applied, a stand is thinned, a ditch is blocked, a crossing structure is built. Whether those actions help depends on a chain of reasoning that must be made explicit. Section 3.1 sets out how habitat quality is evaluated, from the deliberately simple habitat suitability index to resource selection functions fitted to telemetry data. Section 3.2 turns to the manipulation of vegetation through fire, invasive plant control and silviculture. Section 3.3 addresses restoration, where the objective is not to maintain a system but to return one that has been lost. Section 3.4 asks how all of this changes when the climate itself is moving, which converts habitat management from a problem of maintaining a state into a problem of managing a trajectory.

3.1 Habitat Evaluation Methods

3.1.1 What habitat is: the vocabulary of selection

Habitat is the set of resources and conditions present in an area that produce occupancy by a given organism. Three features of that definition carry practical weight. Habitat is species-specific, so there is no such thing as good habitat in the abstract. It is not a vegetation type: a mature oak stand is a vegetation type, whereas the habitat of a particular bat species within it consists of the roost cavities, the flight corridors, the insect prey and the drinking water it contains. In addition, habitat is defined by occupancy, which is a demographic outcome, so a site that animals occupy but in which they consistently fail to reproduce is not good habitat regardless of how often they are seen there. This is the source-sink distinction of Section 1.2.5 restated in habitat terms. It is the reason habitat quality should ultimately be validated against survival and reproduction rather than against density alone.

Four related terms are frequently confused. Use is simply where an animal is recorded. Availability is what the animal could have used, which the analyst must define and which is the most error-prone step in any selection analysis. Selection is use that differs from availability. Preference is selection measured when all resources are equally available, which is achievable in an experiment but almost never in the field. A great deal of confusion in the literature dissolves once these four are kept apart.

3.1.2 The hierarchy of selection

Selection is hierarchical. Johnson's four orders remain the standard organising scheme (Figure 3.1, Table 3.1). First-order selection determines the geographic range of the species. Second-order selection places the home range within that range. Third-order selection governs the use of patches within the home range. Fourth-order selection concerns the choice of individual items or sites within a patch, such as which prey to take or which cavity to nest in.




Figure 3.1. Johnson's four orders of habitat selection. Each panel is an enlargement of the dashed rectangle in the panel to its left. The drivers of selection change with order, so a management question must be matched to the order at which the relevant process operates.

The practical importance of the hierarchy is that the drivers differ among orders, so a study conducted at the wrong scale will find nothing. Climate and biogeographic history dominate the first order; landscape composition and conspecific density dominate the second; forage, cover and predation risk dominate the third; and fine-scale structure dominates the fourth. A proposal to improve habitat by planting shrubs is a third-order or fourth-order intervention, so evaluating it with a first-order climate model is a category error. Equally, a species declining because of first-order climatic change will not be rescued by fourth-order habitat improvement.

Table 3.1. The four orders of habitat selection, with the analyses and management actions appropriate to each.

Order	What is selected	Typical analysis	Management lever
First	The geographic range of the species	Species distribution models over climate and biome layers	Protected area networks; climate adaptation planning
Second	The placement of the home range within the range	Home range centroids against landscape composition	Landscape configuration; reserve siting; connectivity
Third	Use of patches within the home range	Resource selection functions on telemetry relocations	Stand-level treatments: burning, thinning, water provision
Fourth	Items or sites within a patch	Nest-site or prey-item comparisons against paired unused sites	Retention of cavities, snags, cover objects, nest boxes

3.1.3 Habitat suitability index models

The habitat suitability index was developed for the Habitat Evaluation Procedures used in impact assessment. It remains valuable precisely because it is transparent. The analyst identifies the life requisites of the species, normally food, cover, water and reproduction, then selects a measurable variable for each. A suitability index curve is drawn for every variable, mapping the measured value onto a score between zero and one, where zero means the variable alone renders the site unusable and one means it is not limiting (Figure 3.2a to 3.2c). Curves are derived from published relationships, from local data or from structured expert judgement; the basis should always be stated.

The individual scores are then combined into an overall index. The aggregation rule is not a technicality: it embodies a biological assumption, so it changes the answer (Figure 3.2d). An

arithmetic mean assumes that variables compensate for one another, so abundant food can offset scarce water. A geometric mean assumes partial compensation and, because any zero drives the product to zero, correctly forces the index to zero when a requisite is entirely absent. The limiting factor rule takes the minimum, assuming no compensation at all. For requisites that genuinely cannot substitute for one another, such as water and nest cavities, the geometric mean or the minimum is the defensible choice.




Figure 3.2. Constructing a habitat suitability index. (a) to (c) Suitability index curves for three variables, with the score for one surveyed stand marked. (d) The same three scores aggregated three ways, giving indices from 0.42 to 0.65. (e) Habitat units, the product of index and area, compared across four management options. (f) Validation of the index against independently measured density.

Because an index without area is not a management currency, the index is multiplied by the area of the mapped unit to give habitat units (Figure 3.2e). This makes options comparable: a treatment that raises the index modestly across a large area may deliver more than one that transforms a small one. It also exposes the trade in habitat trading and mitigation banking, where losses at one site are offset by gains at another.

The weakness of the approach is that a model built from assumed curves can produce a confident number that has never been tested. Validation against independent field data is therefore essential rather than optional (Figure 3.2f). Published validations are mixed: some models predict density or occupancy well, while others show no relationship at all, usually because the chosen variables were not the limiting ones. An unvalidated habitat suitability index is a hypothesis rather than a measurement, so it should be labelled as such in any assessment document.

3.1.4 Resource selection functions

Where telemetry or survey data are available, selection can be estimated rather than assumed. A resource selection function is any function proportional to the probability of use of a resource unit. In its standard exponential form:

$$w(x) = \exp(\beta_1 x_1 + \beta_2 x_2 + \dots + \beta_k x_k)$$

the coefficients are estimated by comparing the covariate values at used locations against those at a sample of available locations, most often using logistic regression or a Poisson point process formulation (Figure 3.3a and 3.3b). A positive coefficient indicates selection for the covariate; a negative one indicates avoidance.

The design of the availability sample determines what the coefficients mean. Four designs are conventionally distinguished (Table 3.2). The most common analytical error is to define availability carelessly. Drawing available points from the whole study area when the animal could only ever have reached part of it inflates apparent selection; drawing them from a buffer so tight that it contains only used habitat erases it. Availability should be defined by what the animal could plausibly have reached in the relevant time, which is why step selection functions, introduced in Section 2.3.3, are now preferred for fine-scale data: they define availability from the observed movement kernel rather than by an arbitrary polygon.

Table 3.2. Study designs for use-availability analysis, following the standard four-design scheme.

Design	Use measured	Availability measured	Inference supported
I	Population level; individuals not identified	Population level	Population-average selection only
II	Individual level	Population level, shared by all animals	Individual variation in selection from a common landscape
III	Individual level	Individual level, from each home range	Third-order selection within each home range
IV	Individual level, with each use point paired	Individual and time- specific, from the movement kernel	Step selection; fine-scale behavioural inference




Figure 3.3. Resource selection. (a) Used locations compared against a sample of available locations across a covariate surface. (b) The fitted exponential resource selection function, with contours of relative selection strength. (c) The functional response in habitat selection: the estimated coefficient for a resource declines as that resource becomes more available, so a coefficient estimated in one landscape does not transfer unchanged to another.

A further complication is that selection coefficients are not constants of nature. The functional response in habitat selection describes the tendency for animals to select a resource strongly where it is scarce while showing weak selection or indifference where it is abundant (Figure 3.3c). This means that a coefficient estimated in one landscape may not transfer to another with a different composition, which is a common reason that habitat models validated in one region perform poorly in the next.

3.1.5 Species distribution models and their evaluation

At the first and second orders the equivalent tool is the species distribution model, which relates occurrence records to environmental predictors across a region. Presence-absence data support standard regression approaches such as generalised linear and additive models or machine-learning methods such as boosted regression trees and random forests. Presence-only data, which dominate museum and citizen-science archives, are handled by methods such as maximum entropy modelling, which contrasts presences against a background sample. Table 3.3 compares the three families of habitat model treated in this section.

Table 3.3. Three families of habitat model compared.

Approach	Data required	Main strength	Main weakness	Typical order
Habitat suitability index	Expert curves plus mapped habitat variables	Transparent; auditable; needs no animal data	Untested unless validated; expert bias	Third
Resource selection function	Used and available locations	Empirical; quantifies selection strength	Sensitive to the definition of availability	Second and third
Species distribution model	Occurrence records plus environmental layers	Works at broad scales from existing archives	Estimates the realized niche; poor transferability	First and second
Step selection function	Sequential relocations at short intervals	Availability defined by movement; supports connectivity	Requires high-frequency telemetry	Third and fourth
Integrated demographic model	Habitat covariates plus vital rates	Links habitat directly to fitness	Data-hungry; rarely feasible	All

Every habitat model must be evaluated on three distinct properties (Figure 3.4). Discrimination is the ability to rank sites correctly, conventionally summarised by the area under the receiver operating characteristic curve. It is necessary but not sufficient, being inflated when the test data are spatially clustered with the training data, so spatial block cross-validation should be used in place of random splits. Calibration is whether the predicted probabilities match observed frequencies, which discrimination statistics ignore entirely: a model can rank sites perfectly while systematically overpredicting occurrence everywhere. Transferability is whether the model performs where it was not fitted. Table 3.4 lists the pitfalls that most often undermine each.




Figure 3.4. Evaluating a habitat model. (a) Discrimination summarised by the area under the receiver operating characteristic curve. (b) Calibration: whether predicted probabilities match observed frequencies. A model may discriminate well while being badly calibrated. (c) Transferability: when projected conditions fall outside the range of the training data, prediction becomes extrapolation and model behaviour is unconstrained.

Table 3.4. Common pitfalls in habitat modelling and their remedies.

Pitfall	Why it matters	Remedy
Sampling bias in occurrence records	Records cluster near roads and settlements, so the model learns accessibility	Spatial thinning; a background sample with the same bias; effort covariates
Random cross-validation	Spatial autocorrelation inflates apparent accuracy	Spatial block or environmental block cross-validation
Extrapolation into novel conditions	Model behaviour outside the training range is unconstrained	Map novel conditions explicitly; clamp or report where extrapolation occurs
Confusing the realized with the fundamental niche	Biotic constraints are baked into the fitted response	Interpret projections cautiously; use mechanistic models where possible
Ignoring detection	Absences may be non-detections, biasing the fitted response	Occupancy-based distribution models (Section 2.1.4)
No demographic validation	A high-suitability sink is indistinguishable from a source	Validate against survival and productivity, not density alone

3.2 Vegetation and Land Management

3.2.1 Fire regimes and the legacy of suppression

A fire regime is the characteristic pattern of fire in an ecosystem, described by five components: frequency, intensity, severity, seasonality and extent (Table 3.5). Intensity and severity are frequently conflated but are distinct: intensity is the rate of energy release at the flaming front, whereas severity is the resulting ecological effect, principally the proportion of vegetation killed. A fast-moving grass fire can be intense yet leave a woodland canopy intact, while a slow-smouldering ground fire may be low in intensity yet kill root systems and produce severe effects.

Ecosystems occupy characteristic positions in the space defined by return interval and severity (Figure 3.5a). Tallgrass prairie and longleaf pine savanna burned every few years at low severity, so their biota is adapted to frequent surface fire and depends on it. Boreal forest and chaparral burned at intervals of decades to more than a century at high severity, so their biota is adapted to stand replacement followed by long recovery. Applying the wrong regime is as damaging as applying none: frequent low-severity burning in a chaparral system prevents the seed banking that its shrubs depend on.

Table 3.5. The components of a fire regime and their ecological significance.

Component	Definition	Ecological significance	How it is measured
Frequency	Mean interval between fires at a point	Filters the flora for resprouting or seeding strategies	Fire scars; charcoal records; fire atlases
Intensity	Rate of energy release at the flaming front	Determines fire behaviour and suppression difficulty	Flame length; rate of spread; fuel load
Severity	Ecological effect, principally vegetation killed	Determines the successional pathway that follows	Composite burn index; satellite burn ratios
Seasonality	Time of year the fire occurs	Interacts with breeding and with plant phenology	Ignition date records
Extent and pattern	Area burned and the patchiness of the burn	Governs refuge availability and recolonisation	Burn mapping; unburned island area

A century of systematic suppression in fire-adapted systems has produced a characteristic and dangerous outcome (Figure 3.5b). Where historical fires consumed surface fuels every few years, exclusion allows fuel to accumulate continuously while shade-tolerant species establish in the understorey, creating ladder fuels that carry fire into a canopy that historically escaped it. The eventual fire is therefore more severe than any the system experienced historically.




Figure 3.5. Fire regimes. (a) Ecosystems occupy characteristic positions in the space defined by historical fire return interval and fire severity. (b) The suppression legacy: with fire excluded, surface fuel accumulates far beyond the levels the system experienced historically, so the eventual fire burns at a severity outside the range to which the biota is adapted.

3.2.2 Prescribed fire for wildlife

Prescribed burning is the deliberate application of fire under a defined prescription, meaning a specified window of fuel moisture, wind speed, relative humidity and mixing height within which the intended behaviour can be achieved safely. The window is often narrow, so operational capacity rather than ecological knowledge is usually the binding constraint on how much burning is achieved.

For wildlife the important distinction is between the direct and the indirect effects of fire. Direct mortality is usually modest for mobile vertebrates provided the burn is patchy and provided it avoids the breeding season, though it can be severe for slow-moving taxa such as tortoises and for animals in dens or nests. The indirect effect, the change in habitat structure and food supply, is almost always the larger one. Because different species require different successional stages (Figure 3.6b), the management objective is rarely to burn everything to one condition. Patch mosaic burning instead rotates burn units so that early, intermediate and late post-fire stages are present simultaneously (Figure 3.6a). The rationale is the pyrodiversity begets biodiversity hypothesis: a variety of fire histories supports more species than any single history (Figure 3.6c). The evidence is supportive but not universal, since in some systems a single optimal post-fire age supports most species of concern, so pyrodiversity should be treated as a hypothesis to be tested locally rather than as an axiom.




Figure 3.6. Prescribed fire for wildlife. (a) Patch mosaic burning: units are burned on a rotation so that several post-fire stages coexist within one management area. (b) Different guilds peak at different times since fire, so no single post-fire age suits all of them. (c) The pyrodiversity hypothesis: a wider variety of fire histories supports more species.

3.2.3 Invasive plant control

Invasive plants alter habitat structure, displace native food plants, change fire regimes and modify soil chemistry. Control decisions are governed by the invasion curve, which shows how the area occupied grows while the return on each unit of expenditure collapses (Figure 3.7a). Four management phases follow. Prevention is by far the most cost-effective, since a species not introduced needs no control. Early detection and rapid response makes eradication feasible while the infestation is small. Containment limits spread once eradication is no longer achievable. Asset protection defends specific high-value sites once the species is widespread.

Because most agencies face more infestations than they can treat, prioritisation is unavoidable (Figure 3.7b). The three criteria that matter are the impact the species causes, the feasibility of controlling it and the value of what is at risk. A widespread species of modest impact should

not consume the budget that would eradicate a newly arrived transformer species. Table 3.6 compares the principal control methods on efficacy, cost and non-target risk.

Table 3.6. Principal methods of invasive plant control.

Method	How it works	Strengths	Limitations and risks
Manual and mechanical	Pulling, cutting, mowing, grubbing	Precise; no chemical residue; volunteer-friendly	Labour intensive; soil disturbance can favour reinvasion
Chemical, targeted	Cut-stump, basal bark or spot application	High efficacy with limited off-target exposure	Requires trained applicators; resistance can evolve
Chemical, broadcast	Boom or aerial application over an area	Cheap per hectare at large scale	Substantial non-target mortality; drift; public opposition
Biological control	Release of a host-specific natural enemy	Self-sustaining; very low cost per hectare once established	Long approval process; host shifts; suppression is rarely complete
Cultural and ecological	Prescribed fire, grazing, flooding, competitive planting	Addresses the conditions that favour the invader	Slow; must be matched to the invader biology
Integrated management	Sequenced combinations with follow-up	The only approach with a consistent record of success	Requires sustained multi-year commitment




Figure 3.7. Managing invasive plants. (a) The invasion curve: as the invaded area grows, the return on each unit of expenditure falls sharply, which defines four sequential management phases. (b) A prioritisation pathway for a newly detected infestation. (c) Control methods compared on cost and first-year efficacy, with symbol area scaling with non-target risk.

Two cautions temper enthusiasm for control. The first is that removal without follow-up frequently produces secondary invasion, in which a second alien species occupies the cleared ground. Budgets should therefore be planned for the full sequence of treatment, follow-up and revegetation rather than for a single intervention. The second is that an invasive plant may be providing habitat that natives now depend on. Where saltcedar has replaced native riparian woodland along regulated rivers of the south-western United States, it supports nesting by the endangered south-western willow flycatcher, so removal without simultaneous restoration of native vegetation and of the flows that sustain it can leave the bird worse off. Control is a means to an ecological objective, not the objective itself.

3.2.4 Forestry practices for wildlife

Structural complexity is the dominant driver of vertebrate diversity in forests. Silvicultural systems differ chiefly in how much of it they retain (Figure 3.8a to 3.8c, Table 3.7). Even-aged plantation management produces stands of uniform height and age with little deadwood, supporting a species-poor assemblage of generalists. Uneven-aged selection systems maintain a continuous canopy with all age classes present, which suits species requiring closed forest but produces little of the early successional habitat that many declining species need. Retention forestry, in which live trees, snags and downed wood are deliberately left at harvest, is the approach with the strongest evidence base for reconciling timber production with biodiversity: retained structures provide lifeboats through the regeneration phase and legacies that persist into the next rotation.

Three specific structures deserve deliberate management. Snags and cavity trees support a guild of primary excavators and secondary cavity users whose richness rises steeply with snag density and then saturates (Figure 3.8d). Because typical managed stands sit on the steep part of that curve, small increases in retention deliver disproportionate gains. Downed deadwood supports saproxylic invertebrates, fungi, small mammals and amphibians. It is depleted even faster than standing deadwood because it is easily removed for fuel. Large old trees provide cavities, complex bark and heavy limbs that no younger tree can substitute for, so they are irreplaceable on management timescales.

The central tension is rotation length (Figure 3.8e). Mean annual increment peaks well before the age at which old-growth attributes such as large cavities and heavy deadwood begin to accumulate, so an economically optimal rotation systematically prevents those attributes from developing. Extending rotations, retaining permanent old-growth reserves within the managed matrix or accepting a retention level near thirty per cent (Figure 3.8f) are the standard responses. Riparian buffers deserve separate mention because they deliver disproportionate benefit per hectare: they protect aquatic habitat, provide moist microclimate and function as connectivity corridors through an otherwise harvested landscape.




Figure 3.8. Silviculture and wildlife habitat. (a) to (c) Stand structure under even-aged plantation, uneven-aged selection and variable retention systems. (d) Cavity-nester richness rises steeply with snag density before saturating; typical managed stands sit on the steep part of the curve. (e) Mean annual increment peaks long before old-growth attributes develop. (f) Green-tree retention trades timber volume against retained biodiversity.

Table 3.7. Silvicultural systems and their consequences for wildlife habitat.

System	Structure produced	Species favoured	Species disadvantaged
Clearcut without retention	Uniform even-aged regeneration; no legacies	Early successional and open-country specialists	Interior forest species; cavity users; saproxylic taxa
Clearcut with variable retention	Even-aged matrix with aggregated legacies	Both early successional and many forest species	Species needing continuous closed canopy
Shelterwood	Two-aged stand; partial canopy through regeneration	Species tolerating partial canopy	True open-habitat specialists
Single-tree selection	Continuous cover; all age classes	Closed-canopy and cavity-using species	Early successional species; those needing large gaps
Group selection	Continuous cover with small gaps	Gap-dependent species alongside forest interior species	Species requiring extensive early seral habitat
No harvest reserve	Old-growth structure over time	Old-growth dependent taxa; deadwood specialists	Disturbance-dependent species if fire is also excluded

3.3 Restoration Ecology Practices

3.3.1 Principles, references and trajectories

Ecological restoration is the process of assisting the recovery of an ecosystem that has been degraded, damaged or destroyed. The international principles and standards published by the Society for Ecological Restoration provide the framework now most widely used. Two of its elements are worth stating explicitly. The first is the reference ecosystem: a description of the state the site would be in had it not been degraded, adjusted for irreversible changes such as altered climate or extinct species. Without a reference, success cannot be defined. The second is the five-star recovery scale, which scores six attributes from one to five, so progress is reported as a profile rather than as a single number (Figure 3.9b, Table 3.8). A project may score highly on physical conditions while scoring poorly on species composition, which tells the manager precisely where to direct further effort.

Whether intervention is needed at all depends on where the site sits relative to a degradation threshold (Figure 3.9a). If the degrading pressure is removed while the system retains its species pool, its soil and its hydrology, passive recovery may be sufficient and is far cheaper. Beyond a threshold, positive feedbacks hold the system in an alternative stable state: an eroded soil that no longer retains water, a lake locked into a turbid state by sediment resuspension, a grassland converted to an annual-dominated sward that carries fire too frequently for shrubs to establish. Recognising which situation applies is the first analytical task, because active restoration of a system that would have recovered on its own wastes resources, while passive management of a system past a threshold simply fails.




Figure 3.9. Restoration concepts. (a) Trajectories after degradation. Where degradation has not crossed a threshold, passive recovery may suffice; beyond a threshold, the system settles into an alternative stable state from which only active intervention will move it. (b) The five-star recovery wheel, in which six attributes are scored separately so progress is reported as a profile.

Table 3.8. The six recovery attributes used in the five-star assessment, with indicative indicators.

Attribute	What it covers	Example indicators
Absence of threats	Removal of the causes of degradation	Grazing pressure, drainage, pollutant loading, invasive cover
Physical conditions	Substrate and hydrology	Soil profile, compaction, water table depth, hydroperiod
Species composition	Presence of characteristic native biota	Native plant richness; indicator vertebrates and invertebrates
Structural diversity	Vertical and horizontal complexity	Strata present, deadwood volume, patch size distribution
Ecosystem function	Processes operating within the reference range	Productivity, nutrient cycling, recruitment, herbivory
External exchanges	Connection to the wider landscape	Dispersal in and out, gene flow, hydrological connectivity

3.3.2 Wetland re-establishment

Wetlands have been lost more rapidly than any other major habitat type, chiefly through drainage for agriculture. They are among the most frequently restored. The governing principle is that hydrology comes first (Figure 3.10). A wetland is defined by its hydroperiod, meaning the depth, duration, timing and frequency of inundation. Hydroperiod determines the hydric soil conditions, which in turn determine the vegetation zonation and hence the fauna. Planting

wetland vegetation at a site whose hydrology has not been restored is the single most common failure mode in wetland restoration.

The practical sequence therefore begins with diagnosis of why the wetland was lost, since the remedy differs entirely between a drained site, a filled site, an incised channel that has lowered the water table and a site starved by upstream abstraction. Hydrological works follow: plugging ditches, breaking subsurface drainage tile, breaching levees, removing or modifying dams or restoring channel elevation. Only then are soils and microtopography addressed, then vegetation, which in many cases re-establishes from a persistent seed bank without any planting at all. Table 3.9 sets out the main approaches together with their evidence base.




Figure 3.10. Wetland restoration. (a) Hydroperiod determines the zonation of vegetation, which in turn determines the fauna, so the water regime is the controlling variable. (b) The hydrology-first sequence: each step depends on the one before it. Planting before hydrology is restored is the most common cause of failure.

Table 3.9. Approaches to wetland restoration, with their typical outcomes.

Approach	Definition	Typical outcome	Main risk
Restoration	Returning a former wetland to its prior condition	Highest success rate; seed bank often intact	Upstream hydrology may be beyond control
Enhancement	Improving the function of a degraded existing wetland	Reliable gains in specific functions	May trade one function against another
Creation	Constructing a wetland where none existed	Frequently fails to reach reference condition	Wrong soils and hydrology; becomes an open pond
Mitigation banking	Offsetting authorised losses with credits elsewhere	Administratively convenient	Documented net loss of area and function; temporal lag
Dam and levee removal	Restoring channel and floodplain connection	Rapid recovery of riverine wetland processes	Sediment release; legacy contaminants

3.3.3 Rewilding

Rewilding is best understood not as one practice but as a spectrum of approaches united by an emphasis on restoring ecological processes rather than on maintaining a target species composition (Figure 3.11a). At the least interventionist end, passive rewilding removes management and allows succession to proceed. Restoring processes reinstates fire, flooding or large-herbivore grazing. Trophic rewilding reintroduces extirpated keystone consumers in order to recover the interactions described in Section 1.3.3. Taxon substitution uses an ecological analogue where the original species is extinct. Pleistocene rewilding proposes restoring megafaunal function lost thousands of years ago. Intervention, cost and uncertainty all rise along that gradient.

The strongest evidence comes from the trophic end (Figure 3.11b). Beaver reintroduction across Europe and North America reliably produces dam building, raised water tables, buffered peak flows, standing water and deadwood, with measurable increases in amphibians, waterfowl, riparian birds and invertebrates. It equally reliably produces conflict where dams flood roads, culverts or productive land, which is why beaver programmes now build in mitigation from the outset. That combination of large ecological benefit with real and predictable conflict is characteristic of trophic rewilding generally. It is treated further in Chapter 5.

Rewilding attracts sustained criticism on three grounds that a practitioner should be able to answer. The first is the choice of baseline: if ecosystems are always changing, the selection of a target period is a value judgement rather than a scientific finding. The second is welfare, raised most publicly by the Oostvaardersplassen reserve in the Netherlands, where a policy of minimal intervention in a fenced area without predators led to large winter die-offs of introduced herbivores and to a substantial revision of the management regime. The third is that removing management from a cultural landscape may eliminate the semi-natural habitats, such as hay meadows and wood pasture, that centuries of low-intensity use created and that many threatened species now depend on. None of these objections defeats rewilding, but each identifies a decision that must be made explicitly rather than by default.




Figure 3.11. Rewilding. (a) The spectrum of approaches, ordered by the degree of intervention and by the uncertainty of the outcome. (b) A documented trophic rewilding cascade following beaver return to a headwater catchment, including the conflict that reliably accompanies the ecological benefit.

3.3.4 Corridors and stepping stones

Section 1.3.6 established why connectivity matters and Section 2.3.3 covered how it is modelled. Implementation takes three main physical forms (Figure 3.12a to 3.12c, Table 3.10). A continuous corridor is a linear strip of habitat joining two patches, which suits species that

will not cross open ground at all. Stepping stones are a series of small patches close enough for the target species to move between, which is often far cheaper to assemble from existing fragments and works well for mobile taxa. Crossing structures carry animals over or under a linear barrier such as a road or canal. Combined with exclusion fencing they are the standard remedy for road mortality.

Design specifications must be derived from the biology of the target species rather than from convenience. Corridor width is the parameter most often compromised: a strip narrow enough to be entirely edge, in the sense of Section 1.3.5, functions as edge habitat rather than as interior habitat, so it may be used by generalists while remaining impassable for the interior species it was built for. Crossing structure design likewise matters, because preferences differ sharply among groups (Figure 3.12e): large ungulates favour open overpasses with good sightlines, while many carnivores and small mammals readily use underpasses that ungulates avoid.

Evidence for effectiveness should be assessed against a hierarchy (Figure 3.12d). That animals enter a structure is the weakest evidence and the easiest to obtain. That they cross from one side to the other is better. That gene flow between the connected populations has been restored is better still, because it demonstrates that crossings translate into reproduction. That the population growth rate has improved is the strongest evidence and the rarest, since it requires demographic monitoring on both sides over many years. Most published corridor evaluations rest on the lowest two rungs, so claims of demographic benefit should be examined carefully. Corridors also carry costs that deserve honest treatment: they can facilitate the spread of invasive species, pathogens and fire, while concentrating animal movement in ways that can be exploited by poachers.




Figure 3.12. Connectivity in practice. (a) to (c) Three physical solutions: a continuous corridor, stepping stones and a crossing structure with exclusion fencing. (d) The evidence ladder: measures of effectiveness become harder to obtain and more convincing from left to right. (e) Overpasses and underpasses are used differently by different groups, so structure type must be matched to the target species.

Table 3.10. Connectivity solutions and their design considerations.

Solution	Best suited to	Key design parameters	Principal risks
Continuous corridor	Interior specialists that will not cross the matrix	Width sufficient to contain core habitat; vegetation continuity	Too narrow to function; becomes an edge trap
Stepping stones	Mobile species tolerant of short matrix crossings	Inter-patch distance relative to dispersal ability; patch quality	Gaps exceed dispersal capacity for the target species
Riparian network	Many taxa simultaneously; multi-functional	Buffer width; continuity through the catchment	Flood disturbance; agricultural encroachment
Wildlife overpass	Large ungulates; open-country species	Width, substrate, vegetation cover, approach sightlines	High capital cost; poor use if sited without survey data
Underpass and culvert	Carnivores, mesomammals, herptiles	Openness ratio; dry ledge; substrate; guide fencing	Predator ambush; flooding; too narrow for ungulates
Matrix softening	Species that will cross a permeable matrix	Hedgerows, field margins, reduced-intensity land use	Diffuse benefit that is hard to demonstrate

3.4 Climate Change Adaptation

3.4.1 Observed biological responses

Three classes of response to a warming climate are now documented across many taxa (Figure 3.13). The first is spatial: distributions are shifting toward the poles and up in elevation, at median rates on the order of tens of kilometres per decade in latitude and several metres per decade in elevation, with wide variation among taxa and with mobile groups such as butterflies and birds moving fastest. The second is temporal: spring events including budburst, emergence, migration and laying are occurring earlier. The third is the interaction between them. Because interacting species respond at different rates, phenological mismatch develops, in which the peak demand of a consumer no longer coincides with the peak availability of its resource (Figure 3.13b). The classic case is the timing of caterpillar abundance relative to the nestling period of insectivorous birds, where the prey peak has advanced faster than the birds have been able to track it.

These responses have two consequences for habitat management. Species are arriving in protected areas that were designated without them in mind, while others are departing from areas designated specifically for them, so the composition of a reserve is no longer a fixed management target. And because responses are individualistic rather than communal, communities are disassembling and recombining, which means that novel assemblages with no historical analogue are becoming common rather than exceptional.




Figure 3.13. Biological responses to a warming climate. (a) Median rates of range shift in latitude and elevation for major taxonomic groups; values are illustrative of the ranges reported in synthesis studies. (b) Phenological mismatch: the prey peak has advanced faster than the consumer has tracked it. (c) Climate velocity falls steeply where temperature changes rapidly over short distances, so mountainous terrain requires far shorter movements to track a given change than flat terrain does.

3.4.2 Vulnerability, velocity and refugia

Climate change vulnerability is conventionally decomposed into three components (Table 3.11). Exposure is the magnitude of climatic change the species will experience. Sensitivity is the degree to which its fitness responds to that change, which depends on physiological tolerance as described in Section 1.1.2 along with dependence on specific habitats or interactions. Adaptive capacity is the ability to accommodate change through dispersal, plasticity or evolution. A species may be highly exposed yet not vulnerable if its adaptive capacity is high, which is why exposure maps alone are a poor guide to conservation priority.

Table 3.11. Components of climate change vulnerability and the traits that determine them.

Component	What it means	Determining traits	How it is assessed
Exposure	The magnitude and rate of climatic change at the species range	Geographic position; topographic setting	Downscaled climate projections; climate velocity
Sensitivity	The degree to which fitness responds to that change	Thermal tolerance breadth; habitat and diet specialisation; dependence on a single interaction	Physiological experiments; trait databases
Adaptive capacity	The ability to accommodate change	Dispersal ability; phenotypic plasticity; generation time; standing genetic variation	Movement data; common-garden studies; genomics
Realised vulnerability	The combination of the three	All of the above plus barriers to movement	Trait-based and correlative assessment frameworks

Climate velocity is the concept that connects climate to space. It is the rate of change of climate through time divided by the rate of change of climate across space, so its units are distance per time: the speed at which an organism would have to move across the landscape to remain in constant conditions. Its most useful property is that it is low wherever temperature changes rapidly with distance (Figure 3.13c). In mountainous terrain a large change in temperature is available within a short horizontal distance, so velocity is low and a modest movement suffices. Across a flat plain the same change may require hundreds of kilometres of movement. Topographically complex landscapes are therefore disproportionately valuable under climate change, which is a direct and defensible criterion for reserve selection.

The same logic operates at fine scales. Microrefugia are sites whose local conditions are decoupled from the regional climate: cold-air drainage hollows, north-facing slopes in the northern hemisphere, deep canopy shade, boulder fields, springs and caves. Forest canopy alone commonly buffers understorey temperature extremes by several degrees, which can equal or exceed the projected regional warming over decades. Managing for microclimatic buffering, chiefly by maintaining canopy cover, topographic complexity and hydrological function, is one of the few climate adaptation actions that can be taken at the scale of an individual reserve and that delivers benefits regardless of which emissions trajectory is realised.

3.4.3 Choosing a strategy: Resist, Accept, Direct

The most useful practical development in climate adaptation has been the recognition that maintaining the historical state is not always an available option. The Resist-Accept-Direct framework, developed by United States federal agencies and now widely adopted, makes the choice explicit (Figure 3.14a, Table 3.12). To resist is to work against change in order to maintain the historical condition, which is appropriate where the value at stake is high, the change is still reversible and resources are sufficient to sustain the effort indefinitely. To accept is to allow change to proceed without intervening, which is appropriate where change is inevitable or where intervention would be futile. The framework has three properties that make it useful in practice. It legitimises acceptance as a defensible management decision rather than as an admission of failure. It forces the objective to be stated before the action is chosen. And it applies at the level of a management unit rather than a whole agency, so different zones of one reserve may reasonably follow different paths: resisting change in a high-value refugium, accepting it in a large area of low sensitivity while directing succession in a transitional zone.

Table 3.12. The Resist-Accept-Direct framework, with the conditions favouring each option.

Option	What it means	Conditions that favour it	Example
Resist	Actively maintain the historical state	High conservation value; change still reversible; sustained resources available	Shading and augmenting flow in a coldwater stream holding an isolated trout population
Accept	Allow change to proceed without intervening	Change is inevitable; intervention futile or unaffordable; low sensitivity	Allowing a boreal forest margin to convert to aspen parkland after fire
Direct	Steer the system toward a chosen new state	Transformation certain but the destination can be shaped	Planting warm-adapted provenances after a stand-replacing disturbance
Reassess	Revisit the choice as conditions change	Always; none of the three options is permanent	Scheduled review at each management plan cycle

3.4.4 Assisted migration

Where a population cannot reach suitable conditions unaided, deliberate movement becomes a candidate intervention. Three forms are distinguished by the distance involved and by the risk each carries (Table 3.13). Assisted gene flow moves genotypes within the existing range in order to introduce alleles adapted to the conditions the recipient population will experience; it is the least controversial form because it introduces no new species anywhere. Assisted range expansion moves a population just beyond its current margin into habitat it could plausibly have reached unaided given more time. Assisted species migration moves a species far beyond its historical range. It is the form that attracts most objection, because introducing a species to a community with no evolutionary history of it is precisely the mechanism that produces invasions.

A defensible proposal therefore has to clear a sequence of hurdles (Figure 3.14b). It must first show that unaided movement is not achievable within the time available, since protecting and restoring connectivity is preferable wherever it will work. It must show that the source population can sustain the removal. It must assess the risk to the recipient community, including competition, predation, hybridisation and the introduction of pathogens or parasites travelling with the translocated animals. In addition, it should proceed as a monitored trial with agreed criteria for scaling up or stopping, rather than as a single irreversible release. The IUCN guidelines on conservation translocations set out this structure in detail. Chapter 4 develops the operational protocols for translocation generally.

Table 3.13. Forms of assisted migration, with their principal risks.

Form	What is moved, plus its destination	Principal benefit	Principal risk
Assisted gene flow	Individuals or propagules moved within the existing range	Raises adaptive potential without introducing a new species	Outbreeding depression; loss of local adaptation
Assisted range expansion	A population moved just beyond the current range margin	Tracks shifting climate where dispersal is too slow	Ecological surprises in the recipient community
Assisted species migration	A species moved far beyond its historical range	May be the only option for narrowly endemic taxa	Invasion; pathogen co-introduction; competitive displacement
Ex situ holding	Individuals moved to captivity pending suitable habitat	Preserves genetic material when no wild option exists	Adaptation to captivity; high cost; uncertain release prospect

Underlying all of this is a shift in what habitat management is for. Where the objective was once to hold a site in a defined condition, it is increasingly to keep options open: to maintain the topographic and microclimatic diversity that gives species somewhere to go, to keep the landscape permeable enough that they can get there, to retain the genetic variation that adaptation requires, then to monitor closely enough that the choice among resisting, accepting and directing can be revisited as evidence accumulates. Chapter 4 turns to the institutional structures, protected areas and legal instruments through which those choices are actually made.




Figure 3.14. Climate adaptation decisions. (a) The Resist-Accept-Direct framework, with the conditions that favour each option. Different zones of a single reserve may reasonably follow different paths. (b) A decision path for assisted migration, in which each hurdle must be cleared before the next is considered, ending in a monitored trial rather than an irreversible release.

Chapter Summary

- Habitat is species-specific and is defined by occupancy, which is a demographic outcome, so habitat quality should ultimately be validated against survival and reproduction rather than against density.
- Selection is hierarchical. A management question must be matched to the order at which the relevant process operates, since a third-order intervention cannot fix a first-order problem.
- Habitat suitability index models are transparent and auditable, but the aggregation rule encodes a biological assumption about compensation. An unvalidated index is a hypothesis rather than a measurement.
- Resource selection functions estimate selection empirically. The definition of availability drives the result, while the functional response in selection limits how far coefficients transfer between landscapes.
- Habitat models must be evaluated for discrimination, calibration and transferability. Random cross-validation inflates apparent accuracy where data are spatially autocorrelated.
- A fire regime has five components, of which intensity and severity are distinct. Fire exclusion in fire-adapted systems accumulates fuel until the eventual fire exceeds anything the biota is adapted to.
- Patch mosaic burning maintains several post-fire stages at once because different guilds peak at different times since fire. Pyrodiversity is a hypothesis to be tested locally rather than an axiom.
- The invasion curve makes prevention and early response overwhelmingly the most cost-effective phases. Control is a means to an ecological objective, so an invader that has become habitat cannot simply be removed.
- Structural complexity drives forest vertebrate diversity. Retention forestry has the strongest evidence for reconciling harvest with biodiversity, while economically optimal rotations systematically prevent old-growth attributes from developing.

- Restoration requires a reference ecosystem and a recovery profile rather than a single score. Whether passive recovery will work depends on whether a degradation threshold has been crossed.
- Wetland restoration is hydrology first. Planting before the water regime is restored is the commonest failure mode. Created wetlands frequently fail to reach reference condition.
- Rewilding is a spectrum from passive succession to megafaunal restoration, with intervention and uncertainty rising along it. Baseline choice, animal welfare and cultural landscapes are the three standing objections a practitioner must be able to answer.
- Connectivity evidence should be judged on a ladder from entry, through crossing and gene flow, to demographic benefit. Most published evaluations rest on the lowest rungs.
- Climate velocity is low in topographically complex terrain, so such landscapes require shorter movements to track a given change and are disproportionately valuable. Canopy and topography buffer microclimate at manageable scales.
- The Resist-Accept-Direct framework makes the strategic choice explicit and legitimises acceptance as a decision. Assisted migration should clear a sequence of hurdles and proceed as a monitored trial.

Key Terms

Adaptive capacity the ability of a species to accommodate change through dispersal, plasticity or evolution

Assisted gene flow movement of genotypes within the existing range to introduce climate-adapted alleles

Availability the resources an animal could have used, as defined by the analyst

Climate velocity the speed at which an organism must move to remain in constant climatic conditions

Degradation threshold the point beyond which feedbacks hold a system in an alternative stable state

Fire regime the characteristic frequency, intensity, severity, seasonality and extent of fire in a system

Habitat suitability index an index between zero and one aggregating suitability curves for life requisites

Habitat units the product of a suitability index and area, used to compare management options

Hydroperiod the depth, duration, timing and frequency of inundation of a wetland

Microrefugium a site whose local conditions are decoupled from the regional climate

Order of selection the level in Johnson's hierarchy at which a selection decision is made

Patch mosaic burning rotating burn units so that several post-fire stages coexist

Pyrodiversity the variety of fire histories present within a landscape

Reference ecosystem the state a degraded site would be in had it not been degraded, adjusted for irreversible change

Resource selection function a function proportional to the probability of use of a resource unit

Retention forestry harvest that deliberately leaves live trees, snags and downed wood as structural legacies

Trophic rewilding reintroduction of extirpated keystone consumers to restore ecological interactions

Review Questions

1. A habitat suitability index for a woodland bird gives suitability scores of 0.90 for cover, 0.80 for food and 0.05 for nest sites. Calculate the index under the arithmetic mean, the geometric mean and the limiting factor rule, then state which you would report and why.
2. A resource selection function fitted in a landscape with 8 per cent wetland cover gives a strong positive coefficient for wetland. Explain why applying that coefficient in a landscape with 40 per cent wetland cover is likely to mislead.
3. A species distribution model reports an area under the curve of 0.91 on random cross-validation. List three reasons this figure may overstate the model's usefulness for reserve planning.
4. A fire-suppressed ponderosa pine stand has not burned for 90 years. Explain why reintroducing fire immediately may be counterproductive, then outline a sequence of treatments that would restore the historical regime.
5. An agency has funds to treat 200 hectares of invasive plants. It faces a 5 hectare infestation of a newly arrived transformer species and a 2000 hectare infestation of a long-established species of moderate impact. Justify an allocation using the invasion curve.
6. Contrast the wildlife consequences of single-tree selection with those of clearcutting with variable retention. Name one species group favoured by each.
7. A created wetland has been planted with appropriate species but has failed to develop hydric soils or characteristic vegetation after five years. Diagnose the likely cause, then set out the sequence of investigations you would run.

Further Reading

- Manly, B. F. J. et al. *Resource Selection by Animals: Statistical Design and Analysis for Field Studies*. The standard treatment of the designs and models in Section 3.1.4.
- Society for Ecological Restoration. *International Principles and Standards for the Practice of Ecological Restoration*, second edition. Source of the reference ecosystem concept and the five-star recovery scale.
- Lindenmayer, D. B. and Franklin, J. F. *Conserving Forest Biodiversity: A Comprehensive Multiscaled Approach*. Development of the retention forestry material in Section 3.2.4.
- Bond, W. J. and Keeley, J. E. *Papers on fire as a global herbivore and on fire regimes*, for the concepts in Section 3.2.1.
- Schuurman, G. W. et al. *Resist-Accept-Direct: structuring decisions for ecological transformation*, published through the United States Geological Survey and the National Park Service.
- IUCN Species Survival Commission. *Guidelines for Reintroductions and Other Conservation Translocations*. The authoritative framework for the assessment in Section 3.4.4.

CHAPTER IV: PROTECTED AREAS AND WILDLIFE

RESOURCE MANAGEMENT

The three preceding chapters dealt with knowledge: how populations work, how they are measured, how their habitat is evaluated and managed. This chapter deals with institutions. A protected area is a legal instrument as much as an ecological one; a hunting season is a regulation; a recovery plan is an administrative document with statutory force; a disease response is a coordinated act of government. The biology sets the constraints, but outcomes are determined by how rights, revenues and responsibilities are allocated. Section 4.1 treats the design and governance of protected areas on land and at sea. Section 4.2 examines the conditions under which wildlife can be harvested indefinitely, along with the community institutions that make sustainable use work or fail. Section 4.3 covers the recovery of species that are already too depleted for ordinary management. Section 4.4 addresses disease, which respects none of these boundaries and which increasingly forces wildlife management into partnership with human and veterinary health.

4.1 Protected Area Design and Governance

4.1.1 Definition, categories and the global estate

The IUCN defines a protected area as a clearly defined geographical space, recognised, dedicated and managed through legal or other effective means to achieve the long-term conservation of nature with associated ecosystem services and cultural values. Three phrases in that definition carry weight. Clearly defined means the boundary must be mappable and legally established. Long-term means the designation is intended to be permanent rather than a temporary land-use allocation. Conservation of nature must be the primary objective: a site whose main purpose is timber production or watershed protection is not a protected area even if it happens to hold wildlife.

The IUCN assigns each protected area to one of six management categories according to its primary management objective, from strict nature reserves that exclude almost all human presence through to areas managed for the sustainable use of natural resources (Figure 4.1a, Table 4.1). The categories are a description of objectives rather than a ranking of quality. A well-run category VI area that keeps a functioning landscape in community hands may deliver more for wildlife than a category II national park that exists only on paper.

A second designation now sits alongside the formal estate. An other effective area-based conservation measure, abbreviated OECM, is a geographically defined area that is not a protected area yet delivers sustained in-situ conservation of biodiversity as an outcome, whatever its primary purpose. Water catchment reserves, military training lands, sacred groves, permanently closed fishing grounds and many Indigenous territories qualify. The concept matters because it recognises conservation that is already happening under other labels; it also matters because it can be abused, since counting weakly governed land toward an area target inflates the headline figure without changing what happens on the ground. The test is outcome: sustained, effective and documented conservation of biodiversity in place.




Figure 4.1. Protected areas in outline. (a) The six IUCN management categories, ordered by the degree to which human use is restricted. The categories describe objectives, not quality. (b) Global coverage as reported in the Protected Planet Report 2024, against the 30 per cent target of the Kunming-Montreal Global Biodiversity Framework. The third bar shows the far smaller fraction of land that is both protected and well connected.

The global protected area estate now covers approximately 17.6 per cent of land and inland waters together with about 8.4 per cent of the ocean. Target 3 of the Kunming-Montreal Global Biodiversity Framework commits parties to conserving at least 30 per cent of both realms by 2030, which implies a rate of designation far above that achieved since 2020 (Figure 4.1b). Two qualifications matter more than the headline figure. The first is connectivity: only about 8.5 per cent of land is both protected and well connected, so much of the estate consists of isolated fragments of the kind analysed in Section 1.3.6. The second is that Target 3 explicitly recognises other effective area-based conservation measures, meaning areas that deliver sustained biodiversity outcomes without being designated as protected areas.

Table 4.1. The IUCN protected area management categories.

Category	Name	Primary objective	Typical human use
Ia	Strict nature reserve	Protect biodiversity and geodiversity with minimal human disturbance	Research and monitoring only, tightly controlled
Ib	Wilderness area	Protect large unmodified areas retaining natural character	Low-impact non-motorised recreation; traditional use
II	National park	Protect ecosystem-scale processes while permitting visitors	Tourism, education, recreation at managed intensity
III	Natural monument	Protect a specific natural feature	Visitation focused on the feature
IV	Habitat or species management area	Maintain or restore particular species or habitats	Active intervention: mowing, grazing, water control
V	Protected landscape or seascape	Safeguard the values of a landscape shaped by people over time	Continued farming, forestry, settlement, tourism
VI	Protected area with sustainable use	Conserve ecosystems alongside sustainable natural resource use	Regulated harvest, grazing, non-timber products

4.1.2 Governance

Category answers the question of what the area is managed for. Governance answers a separate question: who holds authority and who is accountable. The two axes are independent, so any category may in principle be held under any governance type (Figure 4.2, Table 4.2). Recognising this explicitly was one of the more consequential developments in protected area policy, because it made visible a very large estate of territories conserved by Indigenous peoples and local communities that had previously gone unrecorded in national statistics despite delivering biodiversity outcomes at least as good as state-run areas in many comparisons.

The IUCN governance matrix




Figure 4.2. The IUCN governance matrix. Management category and governance type are independent axes, so the same conservation objective may be pursued under state, shared, private or community authority. Shading indicates how commonly each combination occurs rather than how well it performs.

Table 4.2. Governance types for protected areas, with their characteristic strengths and vulnerabilities.

Governance type	Who decides	Strength	Vulnerability
State	National, regional or local government agency	Legal durability; access to public finance	Political reversal; weak local legitimacy
Shared or collaborative	A formal body combining government with other actors	Combines authority with local knowledge	Slow decisions; unclear accountability
Private	Individual owners, companies, trusts, non-governmental bodies	Rapid action; funding independent of the state	Reversible on sale; objectives may shift
Indigenous peoples and local communities	Customary institutions holding recognised rights	Long time horizons; strong local enforcement	Insecure tenure; external pressure on the resource

4.1.3 Designing a network

Historically, protected areas were placed where land was cheap, remote or unproductive, which produced a global estate biased toward rock and ice and away from the fertile lowlands where biodiversity is concentrated. Systematic conservation planning replaced that opportunism with an explicit procedure (Figure 4.3a). Biodiversity features are mapped across planning units,

quantitative targets are set for each feature, existing protection is credited; new areas are then selected to close the remaining gap at least cost.

Two concepts do most of the work. Complementarity means selecting each new site for what it adds to the existing set rather than for its own richness. The practical gain is large: ranking sites by species richness tends to select several sites holding the same widespread species, whereas complementarity-based selection reaches the same representation target with roughly half as many sites (Figure 4.3b). Irreplaceability measures how much the loss of a site would compromise the ability to meet targets anywhere else (Figure 4.3c). A site holding the only population of a narrow endemic has irreplaceability of one, so no amount of protection elsewhere can substitute for it. Together the two concepts convert reserve selection from an argument about which places are nicest into a solvable optimisation problem. Software such as Marxan and its successors implements exactly that.




Figure 4.3. Systematic conservation planning. (a) The standard six-stage procedure. (b) Selecting sites for what they add, rather than by richness ranking, reaches the representation target with far fewer reserves. (c) An irreplaceability surface: darker planning units contribute features that cannot be represented elsewhere.

Design criteria are conventionally summarised as comprehensiveness, adequacy, representativeness and efficiency. Comprehensiveness asks whether all features are included; adequacy asks whether enough of each is protected for persistence; representativeness asks whether the sample captures the variation within each feature; efficiency asks whether the target is met at acceptable cost. Adequacy is the criterion most often neglected, because it requires the viability analysis of Section 2.4 rather than a simple area percentage.

Two further considerations separate a plan from a durable outcome. The first is cost. Land price, opportunity cost and the recurrent cost of management vary by orders of magnitude across a planning region, so a plan that ignores cost buys far less biodiversity per unit of budget than one that treats cost as an explicit objective. Cost-effective selection routinely achieves the same representation target for a fraction of the expenditure required by an area-only rule. The second is permanence. Protected area downgrading, downsizing and degazettement, abbreviated PADDD, has affected a substantial area of the global estate, most often to accommodate extractive industry or infrastructure. Because a plan assumes that existing protection is secure, widespread PADDD silently invalidates the credited baseline. Management effectiveness assessment, using instruments such as the management effectiveness tracking tool, closes the same gap from the other direction by asking whether designated areas are actually staffed, funded and enforced.

4.1.4 Zoning, size and shape

Few protected areas can exclude all human activity, so most are zoned. The model formalised in the UNESCO biosphere reserve programme uses three concentric zones (Figure 4.4a): a strictly protected core, a buffer zone where low-impact activities compatible with the core are permitted; and a transition area where settlement and sustainable enterprise take place. Zoning reconciles conservation with use, but it only works when the buffer genuinely buffers. A buffer that permits activities incompatible with the core simply relocates the boundary inward.

The consequences of size and shape follow directly from the edge analysis of Section 1.3.5. Because edge penetrates a fixed distance while core area scales with the square of linear dimension, small reserves may retain no core at all; elongate reserves retain far less core than compact ones of identical area (Figure 4.4b and 4.4c). This is the practical reason that reserve boundaries drawn along administrative or cadastral lines frequently perform worse than their area suggests.




Figure 4.4. Reserve design. (a) The three-zone biosphere reserve model. (b) Two reserves of identical area and identical edge depth of influence: the compact one retains twice the core of the elongate one. (c) The proportion of a circular reserve that remains core, as a function of area and of the depth of edge influence.

4.1.5 Marine protected areas

Marine systems differ from terrestrial ones in three ways that change reserve design fundamentally. Most marine organisms have a dispersive larval phase, so populations are demographically open and a reserve exports recruits far beyond its boundary. Boundaries are invisible and enforcement is expensive. And the dominant extractive pressure, fishing, removes biomass selectively by size, which alters the reproductive output of the population disproportionately because fecundity in most fishes scales steeply with body length.

A no-take reserve therefore works through a distinctive mechanism (Figure 4.5a). Protection allows individuals to grow large and old, which raises reproductive output per unit biomass by more than the increase in biomass alone. Those large individuals then export benefit in two ways: adults move across the boundary, which is spillover; and larvae are transported by currents to settle in fished areas, which is larval export. Empirically, biomass inside well-enforced no-take reserves typically rises several-fold over one to two decades, with the largest and most commercially valuable species showing the greatest response (Figure 4.5b).




Figure 4.5. Marine reserves. (a) A no-take reserve protects large, old, highly fecund individuals, which export benefit through adult spillover and larval transport. (b) The typical trajectory of fish biomass inside a well-enforced reserve compared with the adjacent fished area. (c) Reserve size involves a trade-off: protection within rises with size, while export per unit area closed peaks at an intermediate size relative to the adult home range.

Reserve size involves a genuine trade-off (Figure 4.5c). A reserve much smaller than the adult home range of the target species protects almost nothing, since animals cross the boundary constantly. A reserve very much larger protects the interior well but exports relatively little per unit of area closed to fishing. Network design guidance therefore recommends multiple reserves, each large relative to adult movement yet spaced within larval dispersal distance of one another, replicating each habitat type and covering a target fraction of each. Table 4.3 sets out the standard criteria. Two failures recur in practice: reserves designated where fishing pressure was already negligible, which yields impressive compliance statistics but no conservation gain; and reserves without enforcement, which are the marine equivalent of a paper park.

Table 4.3. Design criteria for a marine protected area network.

Criterion	Guidance	Reason
Size	Diameter at least twice the adult home range of the focal species	Retains adults through their normal movements
Spacing	Reserves within typical larval dispersal distance of one another	Permits mutual replenishment across the network
Replication	At least three examples of every habitat type	Insures against localised catastrophe
Representation	A defined fraction of every habitat and depth stratum	Captures the full range of communities
Protection level	A substantial no-take core rather than uniform partial protection	Partial protection delivers a small fraction of the benefit
Enforcement	Surveillance capacity funded before designation	Unenforced reserves show no measurable effect

Finally, designation is not permanent in practice. Protected area downgrading, downsizing and degazettement is now documented at scale across many countries, usually driven by extraction, infrastructure or agricultural expansion. Assessments of management effectiveness, of which the Management Effectiveness Tracking Tool is the most widely used, exist precisely because coverage statistics say nothing about whether an area is achieving its objectives. A realistic account of the global estate distinguishes area designated, area effectively managed and area actually delivering outcomes. The three numbers differ substantially.

4.2 Sustainable Wildlife Utilization

4.2.1 Surplus production and maximum sustainable yield

Harvest is possible because populations below carrying capacity produce a surplus. Taking the logistic model of Section 1.2.2, the surplus available each year is

$$H = rN (1 - N/K)$$

which is a dome-shaped function of abundance, zero at N of zero and again at K , with a maximum at N equal to $K/2$ (Figure 4.6a). That maximum is the maximum sustainable yield:

$$MSY = rK/4, \text{ achieved at } N = K/2$$

Expressed as effort rather than abundance, yield rises with harvest rate to a peak at the rate denoted F_{MSY} , then falls as the stock is drawn down faster than it can replace itself (Figure 4.6b). Any yield below the maximum is obtainable at two different stock sizes, one large and one small; the small one is far more vulnerable, which is a recurring source of confusion in fisheries and harvest debates.

MSY is indispensable as a concept yet dangerous as a target, for four reasons. It assumes r and K are known, whereas both are estimated with wide uncertainty; setting a quota from an optimistic estimate of r produces a harvest that exceeds production at every stock size, so the population declines monotonically to collapse (Figure 4.6c). It assumes a constant environment, whereas r varies among years, so a quota that is sustainable on average is unsustainable in poor years. It targets exactly the stock size at which the population is least buffered, since $K/2$ sits at the point of maximum sensitivity. And it treats the population as a single undifferentiated pool, ignoring the age and sex structure that determines its actual productivity. Modern harvest policy therefore treats MSY as an upper limit that should not be approached rather than as a goal.

Two refinements follow from those limitations. The first concerns behaviour at low abundance. The logistic model assumes per capita growth is highest when the population is smallest, whereas many species show depensation, meaning that growth falls away at low density through mate limitation, loss of cooperative behaviour or disproportionate predation. Where depensation operates there is a threshold below which recovery does not occur even if harvest stops entirely, so any rule that permits the stock to be drawn down toward that region is unsafe regardless of its computed yield. The second refinement is structural. Age-structured and length-structured assessments replace the single-pool model with explicit tracking of recruitment, growth and maturity, which allows productivity to be judged from the proportion of reproductive potential remaining rather than from biomass alone. Contemporary practice therefore expresses policy as reference points: a target that management aims for, a limit that must not be crossed and a trigger that forces a management response, with the distance between them scaled to the uncertainty in the assessment.




Figure 4.6. Maximum sustainable yield. (a) The surplus production curve from a logistic model, with MSY at half of carrying capacity. (b) Equilibrium yield against harvest rate, with the peak at F_{MSY} and overfishing beyond it. (c) The fragility of a fixed quota: if the intrinsic rate of increase is lower than assumed, the quota exceeds production at every stock size, so the population declines to collapse.

4.2.2 Harvest control rules

A harvest control rule specifies how the permitted take depends on the estimated state of the population. The choice of rule matters more than the choice of target (Figure 4.7a, Table 4.4). A fixed quota removes the same number every year regardless of abundance, so the proportion taken rises as the population falls, which is destabilising in precisely the circumstances that matter. A fixed proportion removes a constant fraction, so the take falls automatically as abundance falls, which is self-correcting. A threshold rule sets take to zero below a specified escapement level, then increases it with abundance above that level, which protects the stock at low abundance while still permitting substantial yield when it is healthy.

Simulation under environmental variability makes the difference concrete (Figure 4.7b and 4.7c). A fixed quota set at the estimated MSY drives the median population down and carries a high probability of quasi-extinction. A fixed proportion holds the population at a lower but stable level with good yield. A threshold rule sacrifices a little yield in good years for a substantially lower probability of collapse. Since fixed quotas remain administratively attractive because they are simple to publish and enforce, this is a case where the easiest rule to implement is the worst one to adopt.




Figure 4.7. Harvest control rules compared. (a) Permitted harvest as a function of population size under three rules. (b) Median trajectories from 400 simulated populations with environmental variation. The fixed quota collapses the population; the threshold rule holds it highest. (c) The resulting trade-off between the probability of quasi-extinction and mean annual yield.

Table 4.4. Harvest control rules and their properties.

Rule	How it works	Behaviour as abundance falls	When it is appropriate
Fixed quota	A constant number is taken each year	The proportion taken rises: destabilising	Only where abundance is estimated precisely each year
Fixed proportion	A constant fraction of the estimate is taken	The take falls proportionally: self-correcting	Where a reliable annual index exists
Threshold or escapement	No take below a floor; take rises above it	Take goes to zero before the stock is depleted	Precautionary default for uncertain systems
Fixed effort	Effort is capped rather than the catch	Catch falls with abundance, but efficiency creep erodes the cap	Where catch is hard to monitor but effort is visible
Adaptive rule	The rule itself is revised as data accumulate	Depends on the underlying rule	Long-running programmes with formal monitoring

4.2.3 Selective harvest and its consequences

Harvest is rarely random with respect to individuals. The pattern of selectivity determines both the demographic and the evolutionary consequences. Demographically, the elasticity analysis of Section 1.2.4 identifies which class matters: in long-lived species the population growth rate is most sensitive to adult female survival, so a harvest directed at prime-aged females is far more costly per animal than one directed at yearling males (Figure 4.8b). This is the biological justification for the male-biased, age-structured harvests used in most regulated ungulate and carnivore systems.

Evolutionarily, harvest that selects consistently on a heritable trait is a selective agent (Figure 4.8a). Trophy hunting that targets the largest-horned males imposes selection against the very trait being sought. Long-term datasets from intensively hunted populations have documented declines in horn and antler size over decades. Comparable size-selective effects are well documented in commercially fished stocks, where selective removal of large individuals is associated with earlier maturation at smaller size. The management responses are to raise the minimum age at harvest so that males breed before becoming vulnerable, to cap offtake as a fraction of the mature male cohort; or to rotate hunting pressure among areas.

The third consideration is whether harvest mortality is compensatory or additive, as introduced in Section 1.2.3. At low harvest rates in a population well below carrying capacity, much harvest mortality substitutes for natural mortality that would have occurred anyway. Beyond some rate, compensation is exhausted and further harvest is fully additive (Figure 4.8c). Because the transition point is species-specific and rarely known precisely, harvest rates are normally set well inside the compensatory range.




Figure 4.8. The consequences of selective harvest. (a) Sustained selection against large horns produces a measurable decline in the trait over decades. (b) The elasticity of population growth to each class shows where harvest is most costly. (c) Harvest mortality is partly compensatory at low rates and becomes additive beyond some threshold, which is rarely known precisely.

4.2.4 Community-based management

Where wildlife imposes costs on rural people through crop damage, livestock loss and restricted land use, while its benefits accrue to distant states or tourists, the rational local response is to tolerate poaching and to convert habitat. Community-based natural resource management addresses that asymmetry by devolving rights over the resource, so that the people who bear the costs also capture the benefits (Figure 4.9a). Zimbabwe's CAMPFIRE programme and the Namibian communal conservancies are the most studied examples; community forestry programmes across South and Southeast Asia apply the same logic to timber and non-timber products.

The evidence is mixed. The pattern of success is instructive. Programmes work where rights are securely devolved rather than nominally delegated, where the resource generates enough revenue to be worth managing, where the community institution is legitimate and where benefits reach households rather than stopping at an elite. They fail where the state retains the power to revoke rights, where revenue is captured by intermediaries, where the community is not a community in any functional sense; or where the wildlife cannot generate income sufficient to offset the damage it causes. Ostrom's design principles for common-pool resources predict this pattern well (Figure 4.9b, Table 4.5). They remain the most useful diagnostic available when a programme is being designed or when one is failing.




Figure 4.9. Community-based management. (a) The intended causal chain, from devolved rights through revenue to local tolerance of wildlife. Each link is a point at which programmes fail in

practice. (b) The design principles for durable common-pool resource institutions, with the frequency each is met in programmes judged successful.

Table 4.5. Design principles for enduring common-pool resource institutions, applied to community wildlife management.

Principle	What it means for wildlife	Common failure
Clearly defined boundaries	Both the resource and the group entitled to it are defined	Mobile wildlife crosses boundaries that revenue does not
Rules fitted to local conditions	Quotas and seasons match local ecology and livelihoods	Uniform national regulations imposed regardless of context
Collective choice arrangements	Those affected can modify the rules	Rules set centrally with token consultation
Monitoring by accountable agents	Game guards are drawn from and answerable to the community	Monitoring contracted out with no local accountability
Graduated sanctions	Penalties escalate with the severity of the violation	Either no enforcement or immediate criminal prosecution
Conflict resolution mechanisms	Accessible local means to settle disputes	Disputes reach courts that are distant and expensive
Recognised right to organise	The state recognises the community institution in law	Rights granted by policy that can be withdrawn at will
Nested enterprises	Local institutions sit within district and national tiers	Levels compete rather than reinforce one another

4.3 Endangered Species Recovery

4.3.1 Recovery planning and the One Plan Approach

When a population falls below the point at which ordinary management can sustain it, recovery planning takes over. A recovery plan states the biological criteria that would allow the species to be downlisted or delisted, identifies the threats responsible for the decline, specifies the actions needed to remove them, assigns responsibility and cost; and defines the monitoring that will show whether the actions worked. Plans that omit quantitative, measurable criteria are common and are close to useless, because without them there is no defensible point at which recovery can be declared or the effort redirected.

Historically, in situ conservation of wild populations and ex situ conservation in zoos and breeding centres were pursued by different institutions with different objectives and separate population models. The One Plan Approach, developed through the IUCN Species Survival Commission, treats all populations of a species, wild and managed, as a single metapopulation under a single plan (Figure 4.10). The practical consequences are considerable. Founders are drawn from the wild with the genetic structure of the whole species in mind; the captive population is managed explicitly as a source for release rather than as an end in itself; and decisions about when to bring animals into captivity or return them are taken against one shared population model. Table 4.6 sets out the complementary roles.

The One Plan Approach to species recovery




Figure 4.10. The One Plan Approach. Wild and managed populations are treated as one metapopulation under a single plan with shared objectives, so founders and releases move between them according to one population model rather than two institutional agendas.

Table 4.6. Complementary roles of in situ and ex situ conservation under an integrated plan.

Function	In situ	Ex situ
Primary purpose	Maintain the species in its ecological context	Insure against loss and supply animals for release
Genetic management	Managed through connectivity and translocation	Managed through pedigree analysis and mean kinship
Threat abatement	The only place threats can actually be removed	None; threats are avoided rather than addressed
Behavioural competence	Retained naturally through experience	Must be deliberately maintained through husbandry
Cost per individual	Low once habitat is secured	High and permanent while the programme runs
Chief risk	The threat is not removed in time	Adaptation to captivity; the programme outlives its purpose

4.3.2 Captive breeding

A captive breeding programme has one genetic objective: to retain as much of the source population's gene diversity as possible for as long as necessary. The conventional target is to retain 90 per cent of source gene diversity for 100 years, which is demanding and which drives every other design decision.

Founder number sets the ceiling. A sample of F unrelated founders captures a proportion $1 - 1/(2F)$ of the source gene diversity, so twenty unrelated founders capture about 97.5 per cent while ten capture only 95 per cent (Figure 4.11a). Because that ceiling can never be raised later except by adding new founders, the founder decision is the single most consequential one in the programme. Retention thereafter depends on effective population size, since diversity decays by the factor derived in Section 1.4.5 (Figure 4.11b). Meeting a 90 per cent target for a century typically requires an effective size in the low hundreds, which for most vertebrates means a managed population of several hundred animals distributed across institutions.

Pedigree management is what makes this achievable. Every individual is tracked in a studbook. Breeding pairs are chosen to minimise mean kinship, meaning the average relatedness of an individual to the whole managed population. Prioritising animals with low mean kinship equalises founder representation, which retains diversity far more effectively than random pairing. The counterpart risk is adaptation to captivity: traits favoured in a benign, predator-free, provisioned environment are frequently disadvantageous on release; fitness in the wild can decline measurably within a handful of generations (Figure 4.11c). Countermeasures include minimising generations in captivity, maintaining gene flow with wild populations where any remain, providing environments that demand natural behaviours; and conducting pre-release training in predator avoidance and foraging.

Three supporting techniques extend what pedigree management alone can achieve. Genetic rescue introduces individuals from a differentiated population into a small, inbred one, which restores heterozygosity and commonly produces a rapid demographic response; the Florida panther remains the best-documented case. The risk is outbreeding depression, so donor choice is guided by the degree of divergence, the similarity of the environments and the time since separation. Biobanking preserves gametes, embryos and tissue in cryogenic storage, which effectively suspends the loss of diversity for the material held; a founder whose semen is banked can contribute genes decades after its death, which raises effective size without raising the number of animals maintained. Assisted reproductive technologies, including artificial insemination and embryo transfer, allow that stored material to be used across institutions without moving live animals, which removes both the transport risk and much of the disease risk. All three are adjuncts rather than substitutes, since none of them restores a habitat or removes a threat.




Figure 4.11. Genetic management of a captive population. (a) Founder number sets a ceiling on the gene diversity that can ever be captured. (b) Retention over time depends on effective population size, against the conventional target of 90 per cent for 100 years. (c) Fitness on release declines with generations in captivity unless the programme is actively managed against adaptation to captivity.

4.3.3 Reintroduction and other conservation translocations

The IUCN guidelines classify deliberate movements of organisms for conservation benefit into four categories along an axis of increasing ecological risk (Figure 4.12a, Table 4.7). Reinforcement adds individuals to an existing population. Reintroduction re-establishes a population within the indigenous range where the species has been extirpated. Assisted colonisation moves a species outside its indigenous range to avoid extinction under changing conditions, as discussed in Section 3.4.4. Ecological replacement introduces an analogue species to perform a function lost with an extinct one.

The first and most important question for any of these is whether the cause of the original decline has been removed. Releasing animals into habitat where the threat persists converts a conservation programme into a supply line for the same mortality source. This is the single strongest predictor of failure in reviews of translocation outcomes (Figure 4.12b). Other consistent predictors include releasing enough animals over enough years, releasing into demonstrably suitable habitat, using wild-sourced rather than long-captive stock; and monitoring after release for long enough to detect failure while something can still be done about it.

Release technique matters at the margins. A hard release places animals directly into the wild; a soft release holds them in an enclosure at the site, provides supplementary food and shelter during establishment; it generally improves survival for species that must learn a landscape. Every proposal also requires a disease risk analysis, since translocation is one of the principal mechanisms by which pathogens move between populations, a point developed in Section 4.4.




Figure 4.12. Conservation translocation. (a) The four IUCN categories, ordered by increasing ecological risk. Every category requires feasibility assessment, disease risk analysis and agreed criteria for success and for stopping. (b) Factors associated with success in reviews of translocation outcomes, of which removal of the original cause of decline is by far the strongest.

Table 4.7. Categories of conservation translocation, with representative examples.

Category	Definition	Example	Principal risk
Reinforcement	Adding individuals to an existing conspecific population	Genetic rescue of the Florida panther	Outbreeding depression; disease introduction
Reintroduction	Re-establishing a population inside the indigenous range	Black-footed ferret; Arabian oryx; California condor	The original threat has not been removed
Assisted colonisation	Moving a species outside its indigenous range to avoid extinction	Montane species moved poleward or upslope	Invasion of the recipient community
Ecological replacement	Using an analogue to restore a function lost with an extinct species	Aldabra tortoises substituting for extinct Mascarene tortoises	Unanticipated ecological effects

4.4 Disease Ecology and Health Management

4.4.1 The epidemiological essentials

Wildlife disease management rests on a small number of quantities. The basic reproduction number R_0 is the expected number of secondary infections produced by one infected individual introduced into a wholly susceptible population. If R_0 exceeds one the pathogen spreads; if it is below one the outbreak fades (Figure 4.13a). In the simplest compartment model dividing the population into susceptible, infected and recovered classes,

$$R_0 = \beta / \gamma$$

where β is the transmission coefficient and γ the rate at which infected individuals recover or die.

The single most consequential distinction for management is between density-dependent and frequency-dependent transmission (Figure 4.13b). Under density-dependent transmission, contact rate rises with host density, so R_0 rises with density and there exists a threshold density below which the pathogen cannot persist. Reducing host density can therefore break transmission, which is the theoretical basis of culling. Under frequency-dependent transmission, typical of sexually transmitted infections and of pathogens spread through social contact within groups, contact rate is independent of density, so no threshold density exists and culling cannot eliminate the pathogen however far density is reduced.

The same logic yields the critical control threshold. Reducing the susceptible fraction by vaccination or removal lowers the effective reproduction number. Control is achieved once the proportion treated exceeds $1 - 1/R_0$ (Figure 4.13c). For a pathogen with R_0 of two this means half the population; for R_0 of five it means eighty per cent, which is frequently unachievable in free-ranging wildlife and which explains why many control campaigns fail despite substantial effort.




Figure 4.13. The epidemiological essentials. (a) An epidemic in a naive population under a susceptible-infected-recovered model. (b) Under density-dependent transmission there is a threshold host density below which the pathogen cannot persist; under frequency-dependent transmission there is not. (c) The proportion that must be vaccinated or removed to achieve control rises steeply with R_0 .

4.4.2 Drivers and consequences of wildlife disease

Emerging wildlife disease is overwhelmingly a consequence of human activity. Translocation and trade move pathogens between regions that were previously separated, which is how chytrid fungus reached amphibian populations worldwide and how white-nose syndrome reached North American bats from Eurasia. Habitat conversion increases contact between wildlife, livestock and people at the interface. Provisioning and artificial aggregation, whether at feeding stations, water points or supplementary feed sites, raise contact rates far above natural levels. Climate change shifts vector ranges and host phenology. Table 4.8 sets out the major contemporary cases.

Table 4.8. Major wildlife disease events and what each illustrates.

Disease	Hosts affected	Key feature	What it illustrates
White-nose syndrome	Hibernating bats in North America	Introduced fungus; mass mortality at hibernacula	Introduced pathogens meeting naive hosts
Chytridiomycosis	Amphibians worldwide	Two fungal pathogens spread through trade	Trade as a pathway; the value of import controls
Chronic wasting disease	Cervids in North America, Scandinavia, Korea	Prion persisting in soil for years	Environmental reservoirs defeat host-focused control
Highly pathogenic avian influenza	Wild birds; increasingly mammals	Ongoing panzootic of clade 2.3.4.4b with spillover to mammals	Wildlife, livestock and human health are one problem
Devil facial tumour disease	Tasmanian devil	Transmissible cancer spread by biting	Frequency-dependent transmission; no density threshold
Bovine tuberculosis	Badgers, possums, deer, cattle	Multi-host system at the livestock interface	Culling can worsen spread through social perturbation

Disease can drive population declines directly, as with white-nose syndrome. It can also act indirectly by removing a competitor or predator. Two features make wildlife pathogens particularly difficult. Reservoir hosts that tolerate infection can maintain a pathogen indefinitely while a susceptible species is driven toward extinction, so removing the reservoir may be necessary yet is rarely feasible. And pathogens with environmental reservoirs, prions in soil or fungal spores in cave sediment, persist after every host has gone, so recolonisation reinfects.

4.4.3 One Health and the management toolkit

One Health is the recognition that human, animal and environmental health are interdependent, so surveillance, data and decisions should be integrated across the three sectors rather than fragmented among them (Figure 4.14a). Its practical content is institutional: shared laboratory networks, joint outbreak investigation teams, common data standards and pooled budgets. The recent influenza panzootic illustrates the case, since a virus circulating in wild birds moved into poultry, into wild and farmed mammals and into dairy cattle, so that no single agency held either the information or the authority to respond alone.

The management toolkit is summarised in Table 4.9. The choice among tools follows directly from the transmission biology (Figure 4.14b). Surveillance comes first and takes two forms: scanning surveillance, which looks for anything unusual; and targeted surveillance, which tests for a specified pathogen in a specified population. Biosecurity aims to prevent introduction, which is by a wide margin the most cost-effective intervention and which for wildlife means controlling the wildlife trade, disinfecting equipment between sites and screening animals before translocation.

Culling is the most contested tool. It can succeed where transmission is density dependent, where the threshold density is achievable and where reinvasion can be prevented. It fails where transmission is frequency dependent, where an environmental reservoir persists; or where removal disrupts social structure in ways that increase movement and contact. The bovine tuberculosis and badger case is the standard illustration of the last mechanism: culling that disrupted territorial social groups increased ranging behaviour, which spread infection beyond the cull area, a phenomenon termed the perturbation effect. Vaccination avoids that problem entirely and has been strikingly successful where a suitable oral bait vaccine exists, most notably in the elimination of fox rabies from much of western Europe.

Surveillance design deserves the same quantitative care as any other sampling problem. To be confident that a pathogen is absent; or present below some prevalence, requires a sample large enough that failure to detect it is informative. For a large population, the number of animals that must be tested to have a probability $1 - \alpha$ of detecting at least one positive when the true prevalence is p , using a test of sensitivity Se , is


$$n = \ln(\alpha) / \ln(1 - p \cdot Se)$$

so demonstrating freedom from a pathogen at one per cent prevalence with 95 per cent confidence and a perfect test requires about 299 animals; an imperfect test requires proportionally more. The arithmetic explains why claims of disease freedom based on a handful of opportunistic carcasses carry almost no weight. It also explains the appeal of risk-based surveillance, in which sampling is concentrated on the subpopulations, seasons and locations where infection is most likely, since a targeted sample of modest size can be more informative than a large random one.


Table 4.9. Tools for managing disease in free-ranging wildlife.

Tool	How it works	When it works	When it fails
Biosecurity and trade control	Prevents the pathogen arriving at all	Always the cheapest option per unit of risk averted	Once the pathogen is already established
Surveillance	Detects incursion early enough to act	When the sampling design has adequate power	When effort is too low to detect low prevalence
Culling	Reduces host density below the threshold	Density-dependent transmission; closed populations	Frequency-dependent transmission; social perturbation
Vaccination	Reduces the susceptible fraction below the critical level	Where an effective oral or injectable vaccine exists	Where R_0 is high or delivery coverage is low
Movement control	Blocks the pathway between populations	Where transmission depends on human-assisted movement	Where wild movement alone maintains spread
Treatment of individuals	Cures or protects specific animals	Small, accessible, high-value populations	At population scale in free-ranging wildlife
Refuge and resistance management	Protects uninfected subpopulations; supports resistant genotypes	Where isolation can be maintained	Where the pathogen is already ubiquitous

(a) The One Health framework



(b) Deciding how to manage a wildlife disease

**Figure 4.14.** Managing disease. (a) The One Health framework, in which animal, human and environmental health are treated as one interdependent system. (b) A decision pathway that follows from the transmission biology, including the perturbation effect that makes culling counterproductive in socially structured hosts.

Finally, disease risk analysis should be a routine part of every translocation, as noted in Section 4.3.3. The analysis identifies the pathogens plausibly carried by the source population, the pathogens present at the destination to which the released animals will be naive, the consequences of each; and the mitigation available through screening, quarantine, treatment or choice of a different source. Recovery programmes that have skipped this step have repeatedly

introduced pathogens along with the animals, converting a recovery action into a new threat. Chapter 5 turns from disease to the other domain in which wildlife management is inseparable from human interests, namely direct conflict between people and wild animals.

Chapter Summary

- A protected area is defined by a clear boundary, long-term intent and the primacy of nature conservation. The six IUCN categories describe management objectives rather than quality.
- Category and governance are independent axes. Recognising community and private governance made visible a large conserved estate that national statistics had previously ignored.
- Global coverage stands near 17.6 per cent of land and 8.4 per cent of the ocean against a 30 per cent target. Only about 8.5 per cent of land is both protected and well connected.
- Systematic conservation planning replaces opportunism with explicit targets, complementarity and irreplaceability. Adequacy is the criterion most often neglected because it requires viability analysis.
- Zoning reconciles protection with use only when the buffer is genuinely compatible with the core. Small and elongate reserves retain far less core area than their total area suggests.
- Marine reserves work by protecting large, old, fecund individuals that export benefit through spillover and larval transport. Networks need adequate size, spacing, replication and, above all, enforcement.
- Maximum sustainable yield follows from the logistic model as $rK/4$ at half of carrying capacity. It is indispensable as a concept yet dangerous as a target, because it assumes known parameters and a constant environment.
- The harvest control rule matters more than the target. Fixed quotas are destabilising as abundance falls, while threshold rules trade a little yield for a large reduction in the risk of collapse.
- Selective harvest is a selective agent. Sustained trophy hunting can drive measurable declines in the targeted trait; elasticity analysis identifies which class it is least costly to harvest.
- Community-based management works where rights are securely devolved, revenue is sufficient and benefits reach households. Ostrom's design principles predict the pattern of success and failure well.
- The One Plan Approach treats wild and managed populations as one metapopulation. Recovery plans without quantitative criteria cannot say when recovery has been achieved.
- Founder number sets a permanent ceiling on captured gene diversity; effective size determines how fast the remainder is lost. Adaptation to captivity erodes fitness on release within a few generations.
- Removing the original cause of decline is the strongest predictor of translocation success. Every proposal needs a disease risk analysis and agreed criteria for stopping.
- R_0 and the mode of transmission determine what disease management can achieve. Culling can work under density-dependent transmission; under frequency dependence it fails and may backfire through social perturbation.

- One Health integrates surveillance and decision-making across animal, human and environmental sectors, which the current influenza panzootic shows to be a practical necessity rather than a slogan.

Key Terms

Basic reproduction number (R_0) expected secondary infections from one case in a wholly susceptible population

Complementarity selecting each new site for what it adds to the existing reserve set

Critical control threshold the proportion that must be vaccinated or removed, equal to $1 - 1/R_0$

Frequency-dependent transmission contact rate independent of host density, so no threshold density exists

Gene diversity target the convention of retaining 90 per cent of source gene diversity for 100 years

Harvest control rule the specified relationship between estimated abundance and permitted take

Irreplaceability the degree to which losing a site would prevent targets being met elsewhere

Maximum sustainable yield the largest harvest sustainable indefinitely, equal to $rK/4$ in the logistic model

Mean kinship the average relatedness of an individual to the whole managed population

OECD an other effective area-based conservation measure delivering biodiversity outcomes outside formal protection

One Health integration of human, animal and environmental health in surveillance and decision-making

One Plan Approach treating wild and managed populations of a species under a single integrated plan

Perturbation effect increased ranging and contact caused by culling that disrupts social structure

Spillover and larval export the two routes by which a marine reserve delivers benefit beyond its boundary

Threshold density the host density below which a density-dependent pathogen cannot persist

Review Questions

1. A government proposes to meet its 30 per cent commitment by designating a single very large category VI area in a remote region of low productivity. Evaluate the proposal against comprehensiveness, adequacy, representativeness and efficiency.
2. A 400 hectare circular reserve is subject to a depth of edge influence of 200 metres. Calculate the core area and the percentage of the reserve it represents, then state what this implies for an area-sensitive interior species.
3. A fishery estimates r as 0.30 and K as 40,000 tonnes. Calculate MSY and the stock size at which it occurs. State two reasons why setting the annual quota at that value would be imprudent.

4. Contrast a fixed quota with a threshold harvest control rule for a population subject to strong environmental variability. Which would you recommend? What monitoring would it require?
5. Twelve unrelated founders are available for a captive breeding programme. Calculate the maximum proportion of source gene diversity that can be captured, then outline three measures that would slow its subsequent loss.
6. A proposal recommends reintroducing a raptor to a region from which it was extirpated by poisoning 40 years ago. List, in order of priority, the questions you would require answers to before supporting the release.
7. A pathogen with R_0 of four spreads by frequency-dependent transmission in a social carnivore. Explain why a culling programme is unlikely to eliminate it, then identify the tools you would consider instead.
8. A community conservancy generates substantial trophy hunting revenue, but poaching continues to rise. Using Ostrom's design principles, identify the three most likely institutional causes and the evidence you would gather to test each.

Further Reading

- Dudley, N. (editor). Guidelines for Applying Protected Area Management Categories. IUCN. The authoritative source for the material in Section 4.1.1.
- Margules, C. R. and Pressey, R. L. Systematic conservation planning. *Nature*. The foundational statement of the approach in Section 4.1.3.
- Walters, C. J. and Martell, S. J. D. Fisheries Ecology and Management. Full development of surplus production models and harvest control rules.
- Ostrom, E. Governing the Commons: The Evolution of Institutions for Collective Action. The source of the design principles in Section 4.2.4.
- IUCN Species Survival Commission. Guidelines for Reintroductions and Other Conservation Translocations; and the One Plan Approach documentation.
- Wobeser, G. A. Essentials of Disease in Wild Animals, together with the Protected Planet Report for current coverage statistics.

CHAPTER V: HUMAN-WILDLIFE CONFLICT AND COEXISTENCE FRAMEWORKS

Every preceding chapter has treated wildlife management as a technical problem with a technical answer. This chapter treats the case where that assumption fails. When an elephant flattens a season of maize in a single night, when a leopard takes a calf from a boma, when a bear works its way through a suburban rubbish bin, the biology is usually the simplest part of the situation. The difficult part is that people bear a cost they did not choose while the benefits of the animal accrue somewhere else. Section 5.1 examines what drives conflict and why visible damage is only its surface. Section 5.2 covers the physical and technical toolkit together with the evidence on what each tool achieves. Section 5.3 turns to the social dimension, which determines whether any technical measure survives its first year. Section 5.4 examines the economic and policy instruments that redistribute the costs of living alongside wildlife. The organising idea throughout is coexistence: not the absence of conflict, which is unattainable, but its management to a level that people can live with and wildlife can survive.

5.1 Drivers of Conflict

5.1.1 What conflict is

The IUCN Species Survival Commission defines human-wildlife conflict as struggles that arise when the presence or behaviour of wildlife poses an actual or perceived direct and recurring threat to human interests or needs, leading to disagreement between groups of people and negative impacts on people or wildlife. Four elements of that definition repay attention. The threat may be perceived rather than measured. It must be recurring, which separates conflict from a single unlucky encounter. It leads to disagreement between groups of people, meaning that the conflict is ultimately between humans about wildlife rather than between humans and wildlife. And the impacts run in both directions.

That last point is why the field has moved away from framing the animal as the problem. Retaliatory and preventive killing is now recognised as one of the largest sources of mortality for many large vertebrates. A joint assessment by WWF and the United Nations Environment Programme concluded that conflict-related killing affects more than three-quarters of the world's wild cat species, along with many other carnivores and large herbivores. Target 4 of the Kunming-Montreal Global Biodiversity Framework names the issue explicitly, committing parties to manage human-wildlife interactions so as to minimise conflict for coexistence. For many species conflict is now the principal proximate threat.

The most useful analytical move is to separate the layers of a conflict (Figure 5.1a). At the surface lie the visible material impacts: the eaten crop, the killed animal, the injured or dead person. Beneath them lies a layer of experience that is rarely measured: fear, disturbed sleep, the labour of guarding a field at night, children kept home because the path is dangerous, the loss of any sense of control over one's own land. Beneath that lies conflict between people: disputes over who has authority, promises of benefit never delivered, memories of eviction,

distrust of the agency. At the base lie questions of identity and value. Interventions aimed only at the top layer can reduce damage without reducing conflict, which is the commonest disappointment in this field.

Table 5.1. A typology of human-wildlife impacts, with the measurement each requires.

Impact class	Examples	Typically measured as	Usually missed
Direct material	Crops eaten, livestock killed, stores raided, infrastructure damaged	Incidents per household per season; value of loss	Losses too small to report; repeat victims
Direct to people	Injury, death, disease transmission	Case counts held by health or wildlife agencies	Near misses; long-term trauma
Opportunity cost	Land left unplanted, night guarding, longer routes, schooling forgone	Rarely measured at all	Often exceeds the direct loss
Psychological	Fear, anxiety, sleep loss, loss of autonomy	Occasional survey instruments	The strongest driver of intolerance
Impacts on wildlife	Retaliatory killing, preventive killing, removal, habitat clearance	Carcass counts; permits issued	Cryptic killing, which is most of it




Figure 5.1. The structure of human-wildlife conflict. (a) Visible impacts sit above layers of experience, social conflict and value that determine whether an intervention will be accepted. (b) The reinforcing loop through which damage, lost tolerance and retaliation drive one another; every effective programme must break at least one of its links.

The dynamics matter as much as the structure. Damage reduces income and raises fear, which reduces tolerance, which raises the probability of retaliatory killing, which reduces both the wildlife population and the community's stake in its survival (Figure 5.1b). The loop is reinforcing, so a conflict left to run does not settle at an equilibrium. It escalates until either the wildlife population is locally extinguished or the people leave. Effective programmes break at least one link, whether by reducing damage, by raising the benefit of the animal's presence, by responding rapidly enough to preserve trust or by giving the community authority over the decision.

Drivers operate at three levels. Proximate drivers are the mechanics of the encounter: an unguarded field beside cover, an unsecured boma, food waste on a suburban street. Intermediate drivers are the conditions that generate those encounters: land conversion at the protected-area edge, changes in cropping, the recovery of a predator population under legal protection, roads driven through a migration route. Ultimate drivers are structural: population growth, insecure tenure, poverty that removes any buffer against loss, governance that concentrates costs locally while distributing benefits nationally. Interventions aimed at proximate drivers can be quick and cheap. They are also easily undone if the deeper levels are untouched.

Table 5.2. Drivers of human-wildlife conflict at three levels, with the intervention appropriate to each.

Level	Examples	Time to change	Typical intervention
Proximate	Unguarded crop at the forest edge; open boma at night; accessible waste; night driving on a wildlife road	Weeks	Guarding, barriers, deterrents, attractant removal, signage
Intermediate	Cultivation extending to the reserve boundary; predator recovery; new road or canal; loss of natural prey	Years	Land-use zoning, corridor design, prey restoration, crossing structures
Ultimate	Insecure tenure; poverty and absence of insurance; centralised authority over wildlife; population growth	Decades	Rights devolution, benefit sharing, financial instruments, policy reform

5.1.2 Crop damage

Crop damage is the most widespread form of conflict by number of people affected. It is also the most consistently underestimated, because losses are spread thinly across many households and because the compensating labour is invisible in any accounting system. Three empirical regularities organise the subject (Figure 5.2).

The first is a strong distance decay. Damage concentrates within a few hundred metres of the source habitat, with the decay length set by how far the species will travel across open ground. Primates typically raid within about a hundred metres of forest cover; elephants routinely travel more than a kilometre into cultivation in a single night. A small proportion of farms therefore bears most of the loss, which is what makes targeted mitigation efficient. The second regularity is seasonality. Damage tracks crop phenology, peaking as grain fills and ripens, when nutritional value is highest and defensive investment lowest. Guarding must therefore be seasonal to be affordable; an intervention evaluated outside the peak will appear more effective than it is. The third is selectivity. Cultivated grains and fruits are more digestible, more concentrated in energy and less chemically defended than wild forage, which makes a maize field a superior resource rather than a desperate last resort. Unpalatable crops such as chilli and ginger are consistently avoided.




Figure 5.2. Patterns in crop damage. (a) Damage declines with distance from the source habitat at a rate set by the species. (b) Incidents track the phenology of the crop, peaking as grain matures. (c) Selection indices show that crops are chosen, not encountered at random; unpalatable crops are avoided.

Table 5.3. Characteristic crop-damage profiles of the main problem taxa. Values are indicative teaching figures rather than sourced estimates.

Taxon	Timing	Distance from cover	Damage signature	What works best
Elephant	Night, ripening season	Up to several km	Whole field trampled; stores broken into	Early warning, beehive and chilli fences, well-maintained electric fence
Wild pig	Night, year round	Hundreds of metres	Rooting and uprooting; tubers taken	Trench and fence, population control, guarding
Primates	Day, ripening season	Tens of metres	Selective removal of fruit and cobs	Active human guarding; nothing else reliably works
Deer and antelope	Dawn and dusk	Hundreds of metres	Browsing lines; seedling loss	Fencing, repellents, regulated culling
Granivorous birds	Day, grain fill	Kilometres	Grain stripped from heads	Scaring, netting, synchronous planting

A point on measurement deserves emphasis. Reported damage is not measured damage. Households have obvious incentives to report generously where compensation exists; and equally obvious reasons to under-report where reporting is futile or where cultivation is technically illegal. Comparable estimates require a defined sampling frame, damage assessed in the field against a measured crop stand rather than recalled, exposure expressed per unit area planted; and a control for losses from drought, disease and pests, which frequently exceed wildlife losses even in places where wildlife is blamed for everything.

5.1.3 Livestock depredation

Depredation differs from crop damage in that losses are lumpier, more visible and more emotionally charged. A single night can remove a substantial fraction of a household's capital, since livestock function as savings rather than as income. The emotional weight is

disproportionate to the numbers, which has real consequences: policy responses driven by the salience of the loss rather than by its magnitude are common.

The dominant empirical finding is that husbandry, rather than predator density, explains most of the variation in loss (Figure 5.3a). Animals left out at night, unattended, adjacent to cover suffer losses an order of magnitude greater than animals brought into a reinforced enclosure and accompanied by a herder or a guarding dog by day. Predator-proof enclosures are among the most cost-effective interventions in this field, with reported reductions in night-time losses commonly exceeding eighty per cent where the structure is properly built and maintained.

The second finding is that predation is usually a minor component of total livestock mortality (Figure 5.3b). Disease, malnutrition, drought and theft together account for the large majority of deaths in most extensive systems, with predation contributing something on the order of a tenth. This does not make the loss unimportant, because it falls unevenly and because it is the one cause attributable to a decision someone else made.

The third finding is spatial clustering (Figure 5.3c). Depredation concentrates near cover, near water, on steep or broken terrain that gives a predator an approach; and around particular households whose husbandry is weakest. Risk maps built from verified incidents, terrain and land cover routinely show that a small minority of locations produces most of the events, which allows scarce resources to be aimed rather than spread.




Figure 5.3. Livestock depredation. (a) Loss rates by husbandry practice; the enclosure and the herder matter more than the predator. (b) Predation set against the other causes of livestock mortality. (c) Verified losses cluster within a modelled risk surface, which is what makes targeted mitigation efficient.

Table 5.4. Husbandry measures against depredation, with their demands and limits.

Measure	Mechanism	Evidence of effect	Main limitation
Reinforced night enclosure	Removes the nocturnal opportunity entirely	Large and consistent reductions	Capital cost; requires maintenance and correct siting
Herding by day	Human presence deters most predators	Substantial where herders are attentive	Labour cost falls on children in many systems
Livestock guarding dogs	Detection and confrontation	Strong for canids and mid-sized cats	Selection, training and veterinary care needed
Carcass removal	Denies the reward that creates a habitual killer	Reduces repeat events	Requires prompt reporting
Selective removal of a problem individual	Removes a specific habitual killer	Works where the individual is correctly identified	Frequently the wrong animal is taken

5.1.4 Wildlife in settled and suburban landscapes

Conflict is not confined to rural agriculture. Some of the highest incident rates per household occur in suburban landscapes of the industrialised world, where the combination of moderate housing density, gardens, ornamental planting, waste and legal protection produces conditions many species exploit successfully. The pattern along an urban gradient is characteristic (Figure 5.4a): native species richness peaks at intermediate development intensity, while reported conflict incidents peak slightly further along the gradient, where abundant synanthropic species meet dense human populations with low tolerance and little experience of wildlife.

The species that thrive in these settings share a recognisable profile. They are dietary generalists, behaviourally flexible, tolerant of human proximity, capable of using anthropogenic food; and often released from predation and hunting. Deer, raccoons, foxes, coyotes, wild boar, macaques, bears, geese and gulls recur worldwide. Their conflicts are correspondingly distinctive: vehicle collisions, damage to gardens and structures, waste disruption, disease risk through close contact, aggression during breeding or provisioning; and the specific problem of food-conditioned individuals that lose their wariness altogether.

Two management principles have strong empirical support. The first is that attractants drive the problem. Where deliberate or careless provisioning is removed, incident rates fall sharply; where it continues, no amount of aversive conditioning holds (Figure 5.4c). Bear-resistant waste containers, enforcement of feeding bans, secured compost and fruit-tree management are the interventions with the best return in bear-suburb systems. The second is that incidents are strongly seasonal and predictable (Figure 5.4b), peaking during dispersal, rut and the periods of low light that coincide with commuting hours. Predictability makes targeted warning campaigns, temporary speed limits and seasonal signage worthwhile in a way that year-round measures are not.




Figure 5.4. Wildlife in settled landscapes. (a) Species richness and reported conflict both peak along the urban gradient, with conflict displaced toward the built end. (b) Vehicle collisions concentrate in the dispersal and rut season. (c) Removing the attractant, rather than removing the animal, produces the durable reduction.

5.1.5 Shared space, tolerance and risk

Coexistence is possible because wildlife adjusts. The most general adjustment is temporal. Where people are active by day, many large mammals become more nocturnal, compressing their activity into the hours when human encounter rates are lowest (Figure 5.5a). Studies across continents have documented substantial increases in nocturnality with rising human disturbance. The effect is now understood as one of the main mechanisms permitting large carnivores to persist in densely settled landscapes. The adjustment is not free: it can reduce foraging efficiency and increase energetic costs. It also means a species can be present at high density while remaining almost invisible, which routinely surprises communities when camera-trap results are shared with them.

The human side of coexistence is tolerance, which is not a simple function of damage (Figure 5.5b). Two households experiencing identical losses can hold opposite views; the difference is explained by three factors. The first is benefit: whether the household derives any income, employment or service from the presence of wildlife. The second is trust: whether the responsible authority answers a report promptly, treats the complainant with respect and does what it said it would do. The third is agency: whether the household has any legal means of protecting itself or any voice in the decision. Programmes that reduce damage while ignoring all three frequently find that stated tolerance does not move.

Risk perception follows its own logic, which differs systematically from measured risk (Figure 5.5c). Hazards that are dreaded, involuntary, unfamiliar and vivid are perceived as far more dangerous than their statistics warrant, while familiar and voluntary hazards are discounted. A single fatal attack by a large carnivore can therefore transform the politics of a district while a far larger annual toll from road traffic or snakebite passes without comment. The appropriate response is not to correct people with statistics, which reliably fails and is heard as dismissal. It is to acknowledge the fear as legitimate, to address the features that make the hazard dreaded, notably its involuntary character; and to give people practical means of protection they can operate themselves.




Figure 5.5. The mechanics of shared space. (a) Wildlife shifts activity into the night as human disturbance rises, which allows coexistence at high human density. (b) Tolerance depends on benefit and trust as much as on damage. (c) Perceived danger diverges systematically from measured probability of harm.

A workable quantitative framing treats conflict risk as the product of three components, following standard practice in hazard analysis:

$$\text{Risk} = \text{Hazard} \times \text{Exposure} \times \text{Vulnerability}$$

Hazard is the probability that a damaging animal is present. Exposure is the extent to which people or their assets are in the place where that animal moves. Vulnerability is the degree of loss suffered given an encounter, which is where husbandry, storage, insurance and household wealth enter. The framing is useful because it separates three levers that call for different actions: population and habitat management works on hazard, spatial planning and timing work on exposure; and barriers, husbandry and financial instruments work on vulnerability. It also makes clear why exclusive attention to the animal is usually the least efficient of the three routes.

5.2 Technical and Physical Mitigation

5.2.1 Early-warning and detection systems

An early-warning system converts the presence of an animal into a timely human response. Its architecture has six components (Figure 5.6a): detection, classification, decision against a

threshold, alert, response; and a feedback loop that logs the outcome of every alert so that thresholds can be tuned. The last component is the one most often omitted, which is why so many pilot systems perform well in their trial year and are abandoned in their third.

Detection technologies span a wide range of cost and capability (Table 5.5). Collar-based geofencing sends an alert when a tracked animal crosses a virtual boundary, which is reliable for the collared individual and useless for every animal that is not collared. Camera traps with on-board classification detect any animal passing a fixed point and can run species recognition on the device itself, transmitting an alert rather than an image, which cuts power and bandwidth costs sharply. Acoustic sensors detect species with distinctive calls; infrasonic detection has been applied to elephants.

Table 5.5. Detection technologies for early warning. Cost bands are indicative and change rapidly.

Technology	What it detects	Effective range	Relative cost	Main failure mode
GPS collar geofence	One tracked individual crossing a boundary	Unlimited where network exists	High per animal	Uncollared animals are invisible
Camera trap with on-board model	Any animal at a fixed point	10 to 30 m	Moderate	Fixed points are easily bypassed
Thermal camera	Large bodies at night	Hundreds of metres	High	Cost; power; skilled operation
Acoustic and infrasonic sensor	Vocalising species	Hundreds of metres to km	Moderate	Noise; species overlap
Community observer network	Anything a person sees	Whole landscape	Low in cash, high in time	Fatigue; uneven coverage

Every detection system embodies a trade-off between the probability of detecting a real event and the rate of false alarms (Figure 5.6b). The trade-off is not a technical detail to be settled by engineers. A system tuned for near-certain detection generates frequent false alarms; a community woken repeatedly for nothing stops responding within weeks, at which point the true detections are wasted as well.




Figure 5.6. Early-warning systems. (a) The six-component architecture, including the feedback loop that most implementations omit. (b) Detection performance: every system trades probability of detection against false alarm rate; the operating point is a social decision rather than a technical one.

5.2.2 Deterrents and the problem of habituation

Deterrents work by making a place aversive. They span acoustic devices, lights, chemical repellents such as capsaicin preparations, biological deterrents such as beehive fences; and human presence in its many forms. Their common weakness is habituation: an animal that repeatedly experiences a stimulus without any consequence ceases to respond to it (Figure 5.7b). Habituation is not a defect of a particular device. It is a general property of learning, so any deterrent regime that does not account for it will decay, usually within a few weeks to a few months.

Three design principles slow that decay. The first is unpredictability: randomising timing, location and modality prevents the animal from forming a reliable expectation. The second is pairing the stimulus with a real consequence, so that the association is maintained rather than extinguished; a deterrent backed by an occasional genuine aversive experience retains its effect far longer than one that is always empty. The third is rotation among deterrents with different sensory channels. A regime combining these three routinely sustains effectiveness over a season where a single fixed device fails within a month.

The best-documented biological deterrent is the beehive fence, in which occupied hives are suspended along a wire so that an elephant disturbing the line provokes the colony. A nine-year study across twenty-six farms near Tsavo East National Park in Kenya recorded nearly four thousand elephant approaches. The fences deterred about seventy-six per cent of approaching elephants across the whole study, rising to an average of 86.3 per cent during the six peak crop-growing seasons (Figure 5.7c). The design is attractive because the deterrent is self-maintaining, because it produces honey; and because the farmer owns it. The same study noted its vulnerability to drought, which reduces colony occupancy at exactly the times when crops are most valuable.




Figure 5.7. Deterrence. (a) Cost against effectiveness for the main options; the upper left quadrant is where affordable protection lies. (b) Habituation erodes a fixed deterrent within weeks unless presentation is randomised and paired with a consequence. (c) Beehive fences over nine years at Tsavo East, Kenya.

Two widely used responses deserve scepticism. Translocation of a problem individual is expensive, is frequently followed by homing or by the animal creating the same problem elsewhere, carries disease risk; and often removes an animal that was not responsible. It is appropriate in a narrow set of cases, principally where a specific individual has been reliably identified as a repeat offender. Lethal removal of a habitual individual can be legitimate under a clear legal framework; it can preserve tolerance for the population as a whole by showing that the authority takes the loss seriously. Neither is a population-level solution, because the vacancy created by removing an animal from good habitat is generally filled.

5.2.3 Barriers

Barriers are the most reliable mitigation available where they can be afforded and maintained. A well-built and properly energised electric fence excludes most target species, with the important qualification that both halves of that phrase are load-bearing. The mechanism of an electric fence is psychological rather than physical: it teaches an aversive association, so voltage at the wire must exceed the pain threshold reliably, since a single successful crossing at low voltage teaches the opposite lesson and is extremely difficult to unteach.

The engineering is well understood (Figure 5.8a). Alternating live and earth wires ensure contact on dry ground. Outriggers prevent an animal from pushing against the posts. The lowest wire must deter animals that go under, the height must deter those that go over; gates are the perennial weak point. What determines success in practice is none of these details but the maintenance regime: voltage checked on a schedule, vegetation cleared from the live wires, breaks repaired within days rather than months, gates closed at night; and a named person paid to walk the line. Fences fail through neglect far more often than through design.

The economics of enclosure follow from geometry (Figure 5.8b). Perimeter scales with the square root of area, so the fence cost per hectare protected falls steeply with the size of the block enclosed. A smallholder plot fenced individually is expensive per hectare; the same investment applied to a communally fenced block of many farms protects each hectare at a small fraction of the cost. This is the principal argument for community-scale rather than household-scale fencing; it carries a governance corollary: a shared fence requires a shared institution to maintain it, allocate gate duties and sanction those who leave gates open. The fence is therefore an institutional problem wearing an engineering disguise.

Barriers also have ecological costs. A fence that excludes an elephant from a field also blocks a migration route if it is placed carelessly, fragments populations, causes entanglement mortality; and can convert a permeable landscape into a set of isolated compartments. The design question is therefore not whether to fence but where fencing is compatible with the connectivity requirements developed in Section 3.3. Targeted fencing of high-damage clusters and of specific assets such as grain stores, water points and schools usually delivers more protection per unit of ecological cost than long linear barriers along administrative boundaries.




Figure 5.8. Barriers. (a) The components of a functioning electric fence, with the failure modes that account for most disappointments. (b) Because perimeter scales with the square root of area, fencing cost per hectare falls steeply as the enclosed block grows, which favours community-scale rather than household-scale barriers.

Table 5.6. Comparison of the main physical and technical measures.

Measure	Best evidence of effect	Cost profile	Habituation risk	Key requirement
Electric fence	High exclusion where maintained	High capital, high recurrent	Low if voltage is held	Funded maintenance regime
Beehive fence	About 76% of approaches deterred over nine years	Moderate capital, income-generating	Low	Colony occupancy; drought resilience
Acoustic and light deterrents	Strong initially; decays	Low to moderate	High	Randomisation; pairing with consequence
Guarding by people	Consistently effective while sustained	High in labour, low in cash	Low	Safety; schooling of children
Guarding dogs	Strong for canids and mid-sized cats	Moderate recurrent	Low	Correct breed, rearing and veterinary care
Early-warning alerts	Effective where response capacity exists	Moderate	Not applicable	Acceptable false alarm rate

5.2.4 Spatial planning and targeting

Because damage is spatially clustered, the return on mitigation depends heavily on where it is placed (Figure 5.9c). Where damage is concentrated, treating the worst fifth of sites can prevent something approaching three-quarters of total damage, whereas the same budget spread uniformly prevents only a fifth. The prerequisite is a risk map, built from verified incident records, land cover, terrain, distance to cover and water; and human density, then validated against withheld incidents rather than against the data used to fit it. The methods are the resource selection and species distribution approaches of Section 3.1 applied to a different response variable.

At the landscape scale the design question is the balance between sharing and sparing (Figure 5.9a, 5.9b). Land sharing distributes wildlife-compatible practice across the whole matrix, accepting lower yields in exchange for permeability. Land sparing concentrates wildlife in protected cores while intensifying production elsewhere. Neither dominates in general: sparing performs better for species that tolerate no human activity and need large contiguous areas, while sharing performs better for species that use working landscapes. In conflict management a mixed design is usual, combining a spared core, a buffer where wildlife-compatible land use is supported; and a demarcated zone of settlement where wildlife presence is discouraged. What matters most is that the zones are legible on the ground, agreed with those who live in them; and consistent with the movement requirements of the species.

Linear infrastructure deserves separate treatment. Roads, railways and canals concentrate mortality, sever movement and generate a distinctive class of conflict. The mitigation hierarchy is standard: avoid the alignment where possible, then reduce exposure through fencing that funnels animals toward crossings, then provide crossing structures sized and sited for the target

species, then reduce vehicle speed during the seasons and hours of peak movement. Crossing structures work when they are placed on actual movement routes rather than at administratively convenient points, when their approaches are vegetated appropriately; and when fencing prevents animals from simply walking around them.




Figure 5.9. Spatial approaches. (a) Land sharing spreads wildlife-compatible practice across the matrix. (b) Land sparing concentrates wildlife in a protected core. (c) Because damage is clustered, targeting the worst sites yields far more prevented damage per unit spent than uniform coverage.

5.3 Social Dimensions and Indigenous Knowledge

5.3.1 Why the social dimension decides the outcome

The technical measures of Section 5.2 are necessary. They are almost never sufficient. The reason is visible in the layered structure of Figure 5.1: a fence addresses the top layer while leaving the layers beneath it untouched, so a community that distrusts the agency will not maintain the agency's fence; a household that believes its losses are considered unimportant will not report the incident that would have triggered the response. Evaluations of conflict programmes converge on the finding that the strongest predictor of a durable reduction is not the technology deployed but the quality of the relationship between the people affected and the institution responsible.

Two empirical regularities support this. The first is that trust in the wildlife authority predicts stated tolerance more strongly than experienced damage does (Figure 5.10a). The second is that when benefits are shared, attitudes shift measurably relative to comparison communities; the shift persists (Figure 5.10b). The corollary is uncomfortable for agencies: responsiveness is itself an intervention. A verification visit within twenty-four hours, a payment made when promised, a complaint answered by a named person, all these change behaviour; their absence undoes everything else.

It is equally important to be clear about what communities actually ask for, since it is frequently not what is offered (Figure 5.10c). Across many settings the highest-ranked requests are protection of human life, rapid response to incidents, a say in decisions and fair prompt payment. Cash income from wildlife ranks lower; removal of the animal, which agencies often assume is the demand, ranks lowest of all in many surveys. Programmes designed around what people ask for tend to succeed at lower cost than programmes designed around what institutions find easiest to supply.




Figure 5.10. Social determinants of coexistence. (a) Trust in the responsible authority predicts stated tolerance. (b) A before-after comparison with control villages: attitudes shift when benefits are shared. (c) What affected households rank as important, which differs from what programmes usually offer.

5.3.2 Community-based conservation

Community-based conservation transfers some combination of rights, responsibilities and revenues from the state to a local institution. The logic was developed in Section 4.2.4 for consumptive use; it applies directly to conflict: where a community holds secure rights over wildlife and captures a meaningful share of the value, wildlife becomes an asset whose losses can be offset internally rather than an externally imposed cost.

The evidence is mixed in an instructive way. Programmes succeed where rights are genuinely devolved rather than nominally delegated, where the revenue is material relative to the cost imposed, where the community institution is accountable to its members; and where the state does not exercise an unpredictable veto. They fail where the community is an administrative fiction rather than a group with shared interests, where revenue is captured by intermediaries or elites; and where devolution is reversed as soon as the resource becomes valuable. Ostrom's design principles for common-pool resource institutions remain the most useful diagnostic checklist available.

The distributional question is the crux. Benefits typically arrive as collective goods such as a clinic, a school or a road, while costs fall on individual households as specific losses on specific nights. A household that lost its maize does not experience a share in a new classroom as compensation; the mismatch explains many programmes that appear well designed on paper while failing to shift behaviour. Designs that pair a collective benefit stream with an individual, rapid, loss-linked payment perform better than either alone.

5.3.3 Stakeholder engagement and process design

Engagement begins with identifying who holds a stake and who holds power, which are not the same thing (Figure 5.11a). Affected farmers usually have the highest interest in the outcome and the least influence over the decision, while national ministries often have the reverse. Mapping actors on those two axes clarifies where a process must actively create influence rather than merely invite comment, since a process that consults those who already have power and informs those who do not will reproduce the existing distribution of costs.

It is useful to be honest about the level of participation on offer (Figure 5.11b). The classic ladder runs from manipulation, through information and consultation, to collaboration, partnership and full community control. Each rung is legitimate for some decisions; there is no requirement that every decision sit at the top. What corrodes trust is claiming a high rung while operating on a low one: announcing a partnership and then presenting a completed plan for

endorsement. Stating the rung explicitly at the outset, including what is genuinely open and what is already fixed by law or budget, is a low-cost measure that prevents a large class of subsequent grievance.

Table 5.7. Levels of participation, with what each looks like in conflict management.

Level	What happens	Appropriate when	How it fails
Information	Decisions are explained after they are taken	Statutory duties, emergency action	Presented as consultation
Consultation	Views are gathered; the agency decides	Technical decisions with narrow discretion	Views are gathered and then ignored
Collaboration	The community shapes the options considered	Choosing among mitigation measures	Options are pre-filtered to one
Partnership	Decisions taken jointly; risks and budgets shared	Landscape-scale conflict programmes	Partnership without budget authority
Community control	The community decides and holds the funds	Devolved wildlife management areas	Devolution reversed when value appears




Figure 5.11. Process design. (a) Mapping actors by interest and influence shows where a process must create influence rather than merely invite comment. (b) A ladder of participation; programmes frequently claim the upper rungs while operating on the lower ones.

Conflict between people about wildlife responds to the same techniques used in other disputes. Convening a neutral space, agreeing facts jointly before arguing about values, separating positions from underlying interests; and building agreements incrementally from the areas of overlap are all standard practice in mediation; all are transferable. Where a dispute has become entrenched, a facilitator without a stake in the outcome is usually more effective than the wildlife agency, which is frequently a party to the conflict rather than an arbiter of it.

5.3.4 Traditional and local ecological knowledge

Traditional ecological knowledge is the cumulative body of knowledge, practice and belief about the relationship of living beings to one another and to their environment, held by peoples with long histories of interaction with a place and transmitted across generations. Local ecological knowledge is the broader category held by any long-resident community. Both are of practical value in conflict management for reasons that have nothing to do with sentiment (Figure 5.12a).

The most obvious contribution is temporal depth. Formal monitoring rarely extends beyond a few decades, whereas local knowledge routinely encodes conditions from long before any survey began, which allows a shifted baseline to be detected rather than accepted (Figure 5.12b). The second contribution is spatial resolution: people who walk a landscape daily know which gully the animals use and where a particular individual has been seen, at a grain no practical survey can match. The third is that customary rules of use, where they persist, are already legitimate in the eyes of those expected to follow them. The fourth is that many traditional practices are themselves refined mitigation: rotational grazing away from denning areas, communal herding, sacred groves that function as refuges.

Integration must avoid two symmetrical errors. The first is dismissal, treating local accounts as anecdote to be verified before use. The second is extraction, treating a knowledge system as a database to be mined for data points that are then stripped of context, published without attribution and used to justify decisions the holders never agreed to.

The alternative is co-production: framing the question together, agreeing methods together and owning the findings together. Practical requirements include free, prior and informed consent, agreements on data ownership and use, attribution to knowledge holders, benefit sharing where commercial value arises; and acceptance that some knowledge is not to be recorded at all.

Table 5.8. Principles for working with traditional and local knowledge in conflict management.

Principle	What it requires in practice	What it prevents
Consent before collection	Free, prior and informed consent from the appropriate authority, not merely from individuals	Knowledge taken without sanction
Joint framing	The research question is negotiated rather than presented	Answering a question nobody asked
Attribution	Knowledge holders named as contributors or authors	Uncredited appropriation
Data sovereignty	Agreement on who holds, controls and may release the data	Sensitive site locations exposed
Respect for limits	Acceptance that some knowledge is restricted	Publication of restricted material
Parity of evidence	Both knowledge systems inform the decision; neither validates the other	Local knowledge admitted only when science confirms it




Figure 5.12. Knowledge in conflict management. (a) Co-production treats traditional and scientific knowledge as complementary systems rather than as a hierarchy. (b) Local knowledge frequently records decades of change that began before formal monitoring, which is how a shifted baseline is detected.

5.4 Economic and Policy Solutions

5.4.1 Compensation and insurance

The economic problem of conflict is a mismatch of accounts. The costs of wildlife fall on specific households in specific places, while the benefits accrue to a national or global public. Financial instruments exist to correct that mismatch. The simplest is ex-post compensation: a verified loss is paid for from public funds. It is intuitive, politically attractive and, on the evidence, the weakest instrument in the set.

Its weaknesses are structural rather than incidental (Figure 5.13a). Verification is expensive per claim, because a carcass must be inspected quickly by a qualified person and because attribution is often uncertain. The process leaks claimants at every step: a household that must travel to a district office, complete a form in an unfamiliar language and wait months for part payment will frequently not bother; the households least able to bear the loss are precisely those least able to pursue the claim. Moral hazard is real, since a guaranteed payment weakens the incentive to invest in husbandry. Most seriously, compensation creates no incentive to keep wildlife alive; it pays for damage rather than for outcomes.

Several designs improve on it (Figure 5.13b). Conditional compensation ties eligibility to demonstrated preventive measures, which restores the husbandry incentive. Community-level rather than household-level payment reduces verification costs sharply, since the scheme pays a block based on a verified aggregate rather than adjudicating each case, at the price of leaving distribution to the community. Indexed insurance pays on a measurable proxy such as a predator index or a rainfall index rather than on an inspected loss, which almost eliminates verification cost while introducing basis risk, the possibility that a household suffers a loss the index does not register. Performance payments invert the logic entirely by paying for a conservation outcome rather than for damage.

The Swedish scheme for lynx and wolverine is the standard illustration of performance payment. Rather than compensating Sami reindeer herders for animals killed, the state pays herding communities for verified carnivore reproductions documented on their grazing lands. The incentive is thereby aligned: a herder benefits from carnivores persisting rather than from losses occurring; verification concentrates on a small number of documented reproductions rather than on scattered carcasses. The scheme also shows how such instruments fail. Payment

rates unadjusted for inflation over two decades, disputes about which reproductions are approved; and the concentration of costs on an already marginalised group have eroded trust to the point where the design advantage is largely spent. The underlying relationship, not the payment formula, is what carries a scheme.




Figure 5.13. Financial instruments. (a) The stages at which a compensation scheme loses claimants, each of which also spends trust. (b) Instruments compared on the cost of verifying each case and on whether they create an incentive to keep wildlife alive.

Table 5.9. Financial instruments for redistributing the costs of coexistence.

Instrument	Basis of payment	Verification burden	Principal weakness
Ex-post compensation	Inspected and attributed loss	Very high per claim	No incentive to conserve; leakage of claimants; moral hazard
Conditional compensation	Loss, conditional on preventive measures in place	High	Excludes the poorest, who cannot afford the measures
Community block payment	Aggregate verified loss for a defined group	Moderate	Internal distribution may be inequitable
Indexed insurance	A measurable proxy for expected loss	Low	Basis risk; requires a credible index
Performance payment	Verified conservation outcome, such as documented reproduction	Moderate	Needs indexation and trusted monitoring
Payment for ecosystem services	Delivery of a defined service under contract	Moderate	Fails without conditionality and additionality
Revenue sharing	A defined share of tourism or use revenue	Low	Collective benefit does not offset individual loss

5.4.2 Payments for ecosystem services

A payment for ecosystem services is a voluntary transaction in which a defined service is bought by a service user from a service provider, conditional on the provider securing its supply. Four conditions distinguish a genuine scheme from a subsidy with an environmental label (Figure 5.14a). It must be voluntary on both sides. It must be conditional, so payment stops if the service stops. It must be additional, meaning the service would not have been supplied without the payment. And it must be monitored at a cost proportionate to the value transferred. Failure of conditionality is the commonest fault, because withholding payment from a neighbour is politically hard.

Costa Rica's national programme, established in 1997, remains the most studied example. It pays landholders under five-year and ten-year contracts for forest protection, reforestation, sustainable management and agroforestry, funded principally from a fuel tax and a water charge. Between 1997 and 2019 it enrolled more than 1.3 million hectares under contract, with more than eighteen thousand families participating and roughly 524 million United States dollars invested. Forest cover, which had fallen from over half the national territory in 1950 to about a quarter by 1995, has since recovered substantially. Attribution of that recovery to the programme alone is contested, since a logging ban and structural change in agriculture operated at the same time, which is precisely the additionality problem in practice.

Applied to conflict, payment schemes work best when the service purchased is defined in terms that a household can control. Paying for the presence of a species is problematic, since the household cannot make an animal appear. Paying for habitat retained, for a corridor kept open, for livestock kept out of a calving area during a defined season or for a documented carnivore reproduction are all contractible. The general principle is to buy an action or a verified outcome the provider can influence, at a price that exceeds the opportunity cost of the alternative land use, with a contract long enough to justify the investment it requires.

5.4.3 Ecotourism and revenue models

Wildlife tourism is frequently proposed as the mechanism that makes coexistence pay for itself. It sometimes is. The critical question is not how much revenue the sector generates but how much of it reaches the households bearing the costs (Figure 5.14b). Leakage is severe in most destinations: payments to overseas operators, international carriers, city hotels, imported supplies, national taxation and lodge margins absorb the large majority of a visitor's expenditure, so that the share reaching a community at the point of impact is often in the single digits. A programme premised on tourism revenue without an explicit mechanism for capturing it locally is premised on very little.

The mechanisms that do capture value are well known. Community ownership of the lease or the concession, rather than employment within someone else's, captures margin rather than wage. Local procurement of food, construction and services keeps a further share circulating. Employment quotas matter, though wages alone rarely offset agricultural losses. Concession fees and bed levies paid directly to a community institution provide a predictable base that does not depend on the operator's profitability.

Three cautions apply. Tourism revenue is volatile, as the collapse in visitor numbers during the pandemic of 2020 demonstrated, so a coexistence model that depends on it alone transfers volatility onto households with no reserves. Tourism value is unevenly distributed across species and landscapes: charismatic large mammals in accessible scenic settings attract visitors, while most conflict-prone species in most landscapes attract none. And tourism generates its own conflicts, over access, over disturbance and over who may benefit. A diversified portfolio is more robust than reliance on any single stream.




Figure 5.14. Economic instruments. (a) The four conditions that distinguish a payment for ecosystem services from a subsidy. (b) A representative accounting of where tourist expenditure goes; the share reaching the community bearing the costs is typically small unless capture mechanisms are designed in.

5.4.4 Policy and the coexistence framework

The IUCN Species Survival Commission published the first edition of its guidelines on human-wildlife conflict and coexistence in 2023; they now provide the standard reference framework. Their central argument is that conflict should be approached as a social and political problem in which wildlife is implicated, rather than as a wildlife problem with social side effects. Several of their recommendations amount to a design specification for a national programme.

First, work at the level of the conflict rather than the incident, since a programme organised around individual complaints will never move the underlying drivers. Second, establish who is responsible before an incident occurs, with a named authority, a defined response time and a published procedure. Third, integrate across sectors, since agriculture, roads, land tenure and health all shape conflict while none reports to the wildlife agency. Fourth, embed a financial instrument from the outset rather than adding it after the first crisis. Fifth, monitor the social outcome as well as the biological one, since tolerance and trust predict whether a programme will survive. Sixth, plan for the long term, because relationships take years to build and can be destroyed in a single mishandled incident.

Table 5.10. A design checklist for a coexistence programme, with the diagnostic question for each element.

Element	Diagnostic question	Common failure
Problem definition	Have we described the conflict between people, not only the damage?	A technical fix aimed at the surface layer
Targeting	Is mitigation aimed at the mapped clusters?	Uniform coverage at uniform cost
Response capacity	Can we reach a reported incident within an agreed time?	Alerts with no one to answer them
Participation level	Have we stated which rung of the ladder we are on?	Partnership claimed, consultation delivered
Financial instrument	Does anyone gain from the animal remaining alive?	Compensation only, paid late
Monitoring	Are tolerance and trust measured alongside incidents?	Only carcasses and claims counted
Exit and handover	What continues when the project ends?	Everything stops with the grant

It is worth stating the objective precisely. Coexistence does not mean the absence of conflict, which no landscape containing both people and large wildlife has ever achieved. It means a dynamic condition in which conflict is managed to a level that is tolerable to people and compatible with the persistence of wildlife populations, with institutions capable of absorbing incidents without collapse. Judged against the standard of zero damage, every programme fails. Judged against the standard of a population that persists in a landscape whose people accept its presence, a considerable number succeed.

Chapter Summary

- Human-wildlife conflict is defined by recurring threat, disagreement between groups of people and impacts running in both directions; it is now the principal proximate threat to many large vertebrates.
- Visible damage sits above layers of lived experience, social conflict and value; interventions aimed only at the top layer reduce damage without resolving conflict.
- Damage, lost tolerance and retaliation form a reinforcing loop, so an unmanaged conflict escalates rather than settling.
- Drivers operate at proximate, intermediate and ultimate levels; each level calls for a different intervention on a different timescale.
- Crop damage shows distance decay from cover, tracks crop phenology and is selective by crop, which together make targeted seasonal mitigation efficient.
- Husbandry explains most variation in livestock depredation; predation is usually a small share of total herd mortality but falls unevenly.
- Suburban landscapes generate high incident rates because generalist species meet dense, intolerant human populations; attractant removal is the durable remedy.
- Wildlife accommodates people largely by shifting activity into the night, which permits coexistence at surprisingly high human density.

- Tolerance depends on benefit, trust and agency as much as on damage; perceived danger diverges systematically from measured risk.
- Risk can be decomposed as hazard times exposure times vulnerability, which separates three distinct levers for intervention.
- Early-warning systems trade detection against false alarms; the operating point is a social decision; the feedback loop is what keeps a system alive.
- Deterrents decay through habituation unless presentation is randomised and paired with a real consequence; beehive fences deterred about 76 per cent of elephant approaches over nine years at Tsavo East.
- Barriers are reliable where voltage and maintenance are sustained; because perimeter scales with the square root of area, community-scale fencing is far cheaper per hectare than household-scale.
- Participation should be stated honestly against a ladder that runs from information to community control; traditional knowledge should be co-produced rather than extracted.
- Ex-post compensation is the weakest financial instrument; performance payments, indexed insurance and well-designed payment schemes align incentives better, though all require indexation and trusted verification.

Key Terms

Additionality the requirement that a purchased service would not have been supplied without the payment

Attractant management removal of anthropogenic food and waste that draws wildlife into settlements

Basis risk the chance that an indexed payment fails to match the loss actually suffered

Co-production of knowledge framing questions, agreeing methods and owning findings jointly across knowledge systems

Coexistence a dynamic state in which conflict is managed to a level tolerable to people and compatible with wildlife persistence

Conditionality the rule that payment ceases if the contracted service ceases

Distance decay the decline in damage intensity with distance from the source habitat

Habituation the waning of a response to a repeated stimulus that carries no consequence

Land sharing and sparing two contrasting spatial strategies for allocating production and wildlife across a landscape

Leakage the share of tourist expenditure that never reaches the destination community

Moral hazard the weakening of preventive effort caused by a guarantee against loss

Performance payment payment for a verified conservation outcome rather than for a documented loss

Retaliatory killing killing of wildlife in response to damage already suffered

Tolerance a person's willingness to accept the presence of a species, shaped by benefit, trust and agency

Traditional ecological knowledge cumulative knowledge, practice and belief about local ecological relationships, transmitted across generations

Review Questions

1. A district reports 480 crop-damage incidents in a season across 1,200 farms. Explain why this figure alone cannot support a claim about the severity of the problem, then list the additional quantities you would need.
2. Using the layered model of Figure 5.1, explain why a well-built electric fence funded by an external donor might reduce damage by 80 per cent while leaving stated tolerance unchanged.
3. A predator index shows a stable carnivore population while depredation reports have doubled in three years. Give three explanations that do not require any change in predator numbers.
4. An early-warning system detects 95 per cent of true approaches with 40 false alarms per month. Discuss the trade-off with the community that must respond, then propose how you would choose the operating point with them.
5. Explain why a deterrent that works perfectly in a two-week trial may fail in its third month, then design a regime intended to resist that failure.
6. A community wishes to fence 900 hectares of contiguous farmland rather than 60 separate 15-hectare plots. Estimate the relative fencing cost per hectare of the two options, then state the institutional condition the cheaper option requires.
7. Compare ex-post compensation with a performance payment scheme on verification cost, incentive alignment and equity. State the circumstances under which you would nonetheless choose compensation.
8. A tourism operator states that its lodge generates two million units of revenue annually in a conflict-affected area. What would you need to know before treating that figure as a benefit to the affected households?

Further Reading

- IUCN SSC Human-Wildlife Conflict and Coexistence Specialist Group. IUCN SSC Guidelines on Human-Wildlife Conflict and Coexistence, first edition, 2023. The standard reference framework for this chapter.
- WWF and UNEP. A Future for All: The Need for Human-Wildlife Coexistence, 2021. The global assessment of scale and consequence.
- Woodroffe, R., Thirgood, S. and Rabinowitz, A., editors. People and Wildlife: Conflict or Coexistence? The foundational edited volume for the field.
- King, L. E. and colleagues. Studies on beehive fences at Tsavo East, culminating in the nine-year assessment published in Conservation Science and Practice in 2024.
- Zabel, A. and Holm-Muller, K. Conservation performance payments for carnivore conservation in Sweden, together with the more recent evaluations of how that scheme has fared.
- Wunder, S. Revisiting the concept of payments for environmental services, for the definitional conditions applied in Section 5.4.2.

CHAPTER VI: EMERGING FRONTIERS: TECHNOLOGY, POLICY AND FUTURE OUTLOOK

The five preceding chapters describe a discipline with settled foundations. This one describes its moving edges. Three of them are moving quickly. Sensors and machine learning have changed what can be observed, by orders of magnitude in volume rather than by increments. International law has accumulated into a dense architecture whose weakness is now implementation rather than ambition. Molecular biology has put interventions on the table that were science fiction twenty years ago, some of which are already deployable and most of which are not. Section 6.1 examines what artificial intelligence and remote sensing can and cannot do. Section 6.2 sets out the legal and policy architecture within which all conservation now operates. Section 6.3 treats synthetic biology with the care the subject demands, separating what has been demonstrated from what has been announced. Section 6.4 draws the book together into an account of what an integrated, climate-resilient landscape might look like by 2050 and of the kind of professional the work will require.

6.1 Artificial Intelligence and Remote Sensing

6.1.1 The pipeline, not the model

It is tempting to treat machine learning as the interesting part of conservation technology. It is more useful to treat it as one stage in a pipeline that runs from sensor to decision (Figure 6.1a). Something must sense; the data must be moved and stored in a form that can be found again; a model must propose what is present; a human must verify a sample so that error rates are known; detections must be converted into an estimate of abundance or occupancy; and that estimate must inform a decision with a stated threshold for action. A programme that automates the third stage while neglecting the fifth has bought speed without buying knowledge.

The reason the pipeline matters is arithmetical (Figure 6.1b). Sensor deployment has grown far faster than the human capacity to review what sensors produce. A single camera-trap array can generate millions of images in a season, an autonomous recorder produces thousands of hours of audio; and a satellite constellation images the entire land surface repeatedly. Manual review scales linearly with people, whereas data volume scales with deployment. The gap that opened between the two is the backlog; it is the reason automation moved from curiosity to necessity. Google's SpeciesNet, released as an open-source model in 2025 after years of use inside the Wildlife Insights platform, classifies images into roughly 2,500 categories of mammals, birds and reptiles; the Snapshot Serengeti project used comparable tooling to clear a backlog of eleven million photographs in days rather than years.

6.1.2 Computer vision and its failure modes

Three distinct tasks are routinely conflated. Detection asks whether an animal is present in an image at all, which modern models perform close to perfectly. Classification asks which species it is, which is reliable for common and distinctive taxa and much weaker for cryptic ones.

Individual identification asks which animal it is, which works well where markings are distinctive and remains difficult otherwise. Performance falls sharply along that sequence (Figure 6.2b), so a claim of ninety-nine per cent accuracy should always prompt the question of which task was measured.

Two measures matter more than accuracy. Precision is the proportion of predicted detections that are correct. Recall is the proportion of true events that are detected. A model can achieve high accuracy on an imbalanced dataset by predicting the common class every time, which is why the harmonic mean of the two is the standard summary:

$$F1 = 2 \cdot \text{precision} \cdot \text{recall} / (\text{precision} + \text{recall})$$

Which of the two to favour is a management decision rather than a technical one. A screening system for an invasive species should favour recall, since a missed incursion is far costlier than a false alarm. A system that triggers an armed response should favour precision. The threshold that sets the balance belongs to the manager, not to the model.

Training data requirements follow a predictable curve (Figure 6.2a). A few hundred labelled images per class typically brings a common species to usable accuracy, whereas rare species plateau below it because the tail of the distribution is where the examples are scarcest. This is a structural problem for conservation, since the species of greatest concern are precisely the ones with the fewest images. Transfer learning, synthetic augmentation and few-shot methods mitigate it without removing it.

The most consequential failure mode is domain shift (Figure 6.2c). A model trained on one set of sites learns the backgrounds, the vegetation, the camera models and the lighting of those sites along with the animals. Deployed in a new region it can lose twenty or thirty percentage points of accuracy; the loss is not announced by the model, which continues to report confident predictions. The remedies are routine rather than exotic: hold out entire sites rather than random images when validating, report performance on the held-out sites, retune with a modest number of local labels before deployment; and keep a permanent human verification sample so that drift is detected while it is still small.

A final caution concerns the step most often skipped. A model output is a count of detections, not an estimate of abundance. It must still be corrected for imperfect detection using the methods of Chapter 2:



$$\hat{N} = C / \hat{p}$$

where C is the count and $\hat{p}$ the estimated detection probability. Automation changes how C is obtained. It does nothing whatever to $\hat{p}$; a pipeline that reports raw detections as though they were abundance will produce trends that reflect changes in the sensor network rather than in the population.

Images are only the most visible application. Bioacoustic classifiers now detect the calls of well-studied taxa in continuous recordings with performance comparable to image classifiers, which has made passive acoustic monitoring practical for birds, bats, frogs and cetaceans; and which extends coverage to species that cameras rarely capture. Sequence classifiers assign environmental DNA reads to taxa, with the accuracy of the assignment limited by the completeness of the reference library rather than by the algorithm. Behavioural classifiers infer activity from accelerometry on a tag, which converts a movement record into a budget of foraging, resting and travelling. In every case the same three cautions apply: performance is task-specific, it degrades away from the training distribution; and a classified detection is still a detection rather than a population estimate.

Table 6.1. Machine learning tasks in wildlife science, with their current maturity.

Task	What it does	Maturity	Principal limitation
Animal detection in images	Flags frames containing an animal	Mature	Small or distant animals; heavy vegetation
Species classification	Assigns a species label	Mature for common taxa	Rare species; domain shift between regions
Individual identification	Matches an animal to a known individual	Emerging	Requires distinctive markings and a curated library
Acoustic species recognition	Detects calls in continuous audio	Mature for well-studied taxa	Overlapping calls; background noise
Aerial survey counting	Counts animals in drone or aerial imagery	Emerging	Dense aggregations; double counting
Change detection in imagery	Flags forest loss, burn scars, new roads	Mature	Cloud cover; confusion with natural change
Behaviour classification from tags	Infers behaviour from accelerometry	Emerging	Needs labelled observation to calibrate
Poaching risk prediction	Ranks locations by predicted risk	Experimental	Learns patrol effort rather than poaching

(a) From sensor to decision**(b) Why automation became unavoidable****Figure 6.1.** The technology pipeline. (a) Six stages run from sensor to decision; automation addresses the third, while the fifth is where a detection becomes an estimate. (b) Data volume has grown far faster than human review capacity, which is why automated review became unavoidable.




Figure 6.2. Computer vision in practice. (a) Accuracy against labelled examples per class; rare species plateau below the usable threshold. (b) Precision and recall diverge sharply across detection, classification and individual identification. (c) Domain shift: accuracy falls when a model meets a new region; it recovers after retuning with local labels.

6.1.3 Remote sensing platforms


No platform is best. Each occupies a position in a trade-off between spatial resolution, spatial extent and revisit interval (Figure 6.3a); the correct choice follows from the question rather than from the technology. A study of nest occupancy needs centimetres and can accept a small area. A study of continental land-cover change needs tens of metres and the whole surface. The upper left corner of the resolution-extent plane, where one might image an entire continent at centimetre resolution every day, is empty for reasons of physics and cost rather than engineering ambition.

Drones have become the standard tool at the fine end. They deliver centimetre imagery on demand, they can be flown by a small team; and thermal payloads make them effective for nocturnal counting and for detecting animals under partial canopy. The design of a drone survey is a sampling problem of the kind treated in Chapter 2 (Figure 6.3b): flight lines with defined forward and side overlap, altitude chosen to give the required ground sample distance, ground control for geometric accuracy; and a log of altitude, time and weather so that detectability can be modelled afterwards. Two constraints are frequently underestimated. Disturbance is real, since many species respond to a drone as they would to an aerial predator, so minimum approach altitudes must be established by trial rather than assumed. Airspace regulation is binding in most jurisdictions; permits are not a formality.

At the coarse end, freely available public imagery has transformed what a small organisation can do. Optical missions at ten to thirty metres provide global coverage on a revisit of days, which supports near-real-time detection of forest loss, fire and new roads. Radar penetrates cloud, which matters enormously in the humid tropics where optical imagery is often useless for months. Lidar and radar missions map three-dimensional structure, which relates directly to the habitat variables of Chapter 3. Commercial constellations supply sub-metre imagery where the budget exists, which has enabled direct counting of large animals such as whales, seals and elephants from space in open habitat.

Table 6.2. Remote sensing platforms compared. Figures are indicative and change rapidly.

Platform	Resolution	Extent	Revisit	Best used for
Small drone	1 to 5 cm	Tens to hundreds of ha	On demand	Counts, nests, fine habitat structure
Crewed aircraft	10 to 50 cm	Thousands of ha per day	Chartered	Large-mammal surveys, aerial lidar
Commercial satellite	0.3 to 1 m	Regional	Days, tasked	Direct counting; infrastructure change
Public optical satellite	10 to 30 m	Global	Days	Land cover, deforestation, fire
Radar satellite	5 to 25 m	Global	Days	Cloud-free monitoring; structure; flooding
Spaceborne lidar and radar	25 m and coarser	Global	Mission dependent	Canopy height and biomass

**Figure 6.3.** Remote sensing. (a) Platforms occupy positions in the trade-off between ground sample distance and area covered; the fine-and-broad corner is empty. (b) A drone survey is a sampling design, with overlap, altitude and ground control chosen to make the imagery analysable.

6.1.4 Tracking from space

Animal tracking has been transformed twice: first by the move from very high frequency radio to satellite positioning, then by miniaturisation. The binding constraint has always been mass. The traditional rule of thumb limits a tag to five per cent of body mass; current guidance is stricter, with three per cent treated as an upper bound and one per cent regarded as good practice where achievable (Figure 6.4a). Because tag mass is dominated by the battery; and because battery energy scales with mass, the mass limit sets everything else: how often a position can be taken, whether data can be transmitted rather than stored; and how long the tag will last.

That constraint produces the central design trade-off (Figure 6.4c). Tag life falls close to inversely with fix frequency, so a schedule of one position per day may run for years while a schedule of one position every fifteen minutes may exhaust the same battery in weeks. The choice should follow the question. Migration timing needs a long record at low frequency. Habitat selection at the scale of a home range needs frequent fixes over a season. Fine-scale behaviour needs bursts of high-frequency data separated by long gaps, which is what duty cycling delivers. Solar harvesting relaxes the constraint for species that spend time in sunlight.

Accuracy varies by three orders of magnitude across systems (Figure 6.4b). Doppler-based satellite positioning gives locations accurate to hundreds of metres, which is ample for ocean-basin movement and useless for a woodland home range. Satellite navigation fixes give tens of metres or better. Reverse arrays, in which fixed receivers time a signal from a lightweight transmitter, give sub-metre accuracy within a limited area. The relevant question is never which system is most accurate but whether the error is small relative to the ecological feature being studied.

Table 6.3. Tracking technologies compared.

System	Accuracy	Data return	Typical mass	Main constraint
VHF radio telemetry	Hundreds of metres	Manual, by field team	From under 1 g	Labour; range; no data without an operator
Doppler satellite (Argos)	150 to 1000 m	Remote, global	From about 5 g	Coarse positions; duty cycled
GPS with archival storage	5 to 20 m	Only if the tag is recovered	From about 1 g	Recovery rate governs the sample
GPS with satellite or cell relay	5 to 20 m	Remote, near real time	From about 20 g	Power; network coverage
Reverse GPS array	Under 1 m	Continuous within the array	Under 1 g	Only within the instrumented area
Geolocator (light level)	50 to 200 km	Only if recovered	Under 0.5 g	Very coarse; equinox failure
Accelerometer and biollogger	Not positional	Archival or relayed	Varies	Requires observation to calibrate




Figure 6.4. Tracking constraints. (a) Tag mass is bounded by a proportion of body mass, which limits battery capacity. (b) Positional accuracy across systems spans three orders of magnitude. (c) Because energy is fixed, higher temporal resolution is paid for directly in tag life.

A final component of the pipeline is rarely budgeted for and frequently determines whether any of the rest survives: data management. Standards exist and should be used, since a dataset in a documented schema can be combined with others while a spreadsheet of local conventions cannot. Repositories such as Movebank for tracking data, Wildlife Insights for camera-trap imagery and the Global Biodiversity Information Facility for occurrence records provide durable storage together with an identifier that a paper can cite. Two obligations sit alongside openness. Sensitive location data, above all for species targeted by collectors and poachers,

must be generalised or embargoed before publication, since precise coordinates for a rare species are an operational aid to those who would take it. And data collected in a country or on Indigenous land belong to the people who hold the rights to them, which means agreements on ownership, access and use should be settled before fieldwork rather than after publication.

6.2 Conservation Law, Policy and Treaties

6.2.1 The architecture

Conservation law is layered (Figure 6.5a). Global treaties create obligations for states. Regional instruments translate them into binding rules for groups of states. National law converts obligation into enforceable duty. Subnational rules and customary institutions determine what actually happens on a given hectare. Obligation flows down the layers; capacity, information and legitimacy have to flow up. A treaty commitment that never reaches the lowest layer changes nothing, which is why the implementation gap, rather than the ambition gap, is the dominant problem in this field.

The instruments accumulated over half a century (Figure 6.5b). Ramsar in 1971 protected wetlands, the World Heritage Convention followed in 1972, CITES regulated international trade from 1973; and the Convention on Migratory Species arrived in 1979. The Convention on Biological Diversity in 1992 was the first framework covering biodiversity as a whole, with the Cartagena Protocol on biosafety in 2000 and the Nagoya Protocol on access and benefit sharing in 2010. The Kunming-Montreal Global Biodiversity Framework of 2022 set the current targets, including the thirty per cent protection commitment of Chapter 4 and the coexistence commitment of Chapter 5. The agreement on marine biodiversity beyond national jurisdiction, adopted in 2023, closed the largest remaining spatial gap. The seventeenth Conference of the Parties to the Convention on Biological Diversity is scheduled for Yerevan in late 2026, where progress against the framework will be assessed.

Table 6.4. The principal global instruments, with what each does and how it binds.

Instrument	Scope	Nature of obligation	Characteristic weakness
CITES (1973)	International trade in listed species	Binding; permit system with trade sanctions	Silent on domestic trade and on habitat
CBD (1992)	Biodiversity as a whole, plus access and benefit sharing	Binding in principle; targets are political	No enforcement mechanism; national reporting is uneven
CMS (1979)	Migratory species and their ranges	Binding, with daughter agreements	Depends on all range states joining
Ramsar (1971)	Wetlands of international importance	Listing plus wise-use obligation	Listing does not guarantee management
World Heritage (1972)	Sites of outstanding universal value	Listing; in-danger mechanism	Political pressure against in-danger listing
BBNJ (2023)	Marine biodiversity in the high seas	Binding once in force; area-based tools	Implementation and financing untested




Figure 6.5. The legal architecture. (a) Obligation flows down from global instruments to the field, while capacity and reporting must flow back up. (b) Fifty years of instrument building, from Ramsar in 1971 to the high seas agreement in 2023.

What makes a treaty bite varies. CITES has the strongest mechanism, since the Standing Committee can recommend that Parties suspend trade with a state that fails to implement the convention, which is a commercial sanction with real cost. The Convention on Biological Diversity has no comparable instrument, relying instead on national target setting, periodic reporting and peer pressure, which is why its targets have historically been missed with few consequences. Regional instruments can be stronger than global ones: the European Union's nature legislation is enforceable through the Court of Justice, with financial penalties available against member states. Domestic law in large consumer markets reaches further still. Legislation that makes it an offence to import wildlife or timber taken in violation of another country's law extends the reach of that country's rules into the market that creates the demand, which is often more effective than any measure available at the source.

6.2.2 How CITES works

CITES is worth understanding in detail because it is the treaty with real teeth. It regulates international trade in listed species through three appendices (Figure 6.6a, Table 6.5). Appendix I covers species threatened with extinction, for which commercial international trade is prohibited. Appendix II covers species that may become threatened unless trade is controlled, for which trade is permitted subject to permits. Appendix III covers species that a single Party has listed unilaterally and for which it seeks the cooperation of others.

The scientific heart of the system is the non-detriment finding. Before an export permit for an Appendix II specimen is issued, the exporting state's scientific authority must conclude that the export will not be detrimental to the survival of the species. That is a population-level judgement requiring exactly the material of Chapters 2 and 4: an abundance estimate, a harvest rate, an understanding of productivity, a monitoring system capable of detecting decline. Where those are absent the finding becomes a formality, which is the most common way in which the system fails quietly. A legal acquisition finding is also required, confirming that the specimen was obtained in accordance with national law.

The twentieth Conference of the Parties, held in Samarkand from late November to early December 2025, illustrates the direction of travel (Figure 6.6b). Nearly ninety animal species received new or strengthened listings, including the okapi, several vulture species, Galápagos iguanas, the oceanic whitetip shark, the whale shark and the mobulid rays. Marine taxa and species traded as pets accounted for a large share of the decisions, which reflects where trade pressure has moved. Two decisions in the other direction are equally instructive: the bontebok was removed from the appendices following its recovery; and the Caribbean monk seal was deleted because it is extinct.

Table 6.5. The CITES appendices and what each requires.

Appendix	Criterion for listing	Trade permitted	Documents required
Appendix I	Threatened with extinction and affected by trade	No commercial international trade; non-commercial only	Import and export permits; both authorities must agree
Appendix II	May become threatened unless trade is controlled	Commercial trade permitted if not detrimental	Export permit, resting on a non-detriment finding
Appendix III	Listed unilaterally by one Party seeking cooperation	Trade permitted with documentation	Export permit from the listing state; certificate of origin elsewhere




Figure 6.6. CITES in operation. (a) The decision path from a proposed shipment to a permitted trade, with the non-detriment finding as the scientific step. (b) The pattern of recent listing decisions, which has shifted toward marine taxa and species in the pet trade.

6.2.3 Enforcement

Illegal wildlife trade is a supply chain, so it presents several distinct points at which enforcement can act (Figure 6.7a). At the source, ranger patrols, structured patrol planning and community information networks raise the probability that an offence is detected. In transit, container profiling, detection dogs and forensic tools such as DNA and stable isotope analysis allow a seizure to be linked to a place of origin, which converts an interception into an investigation. At the market, inspection, monitoring of online marketplaces and financial investigation reach the organisers rather than the couriers. Behind all of them lies demand, where behaviour change and social norm interventions operate on a slower timescale.

The economics of deterrence are simple to state. The expected cost of committing an offence is

$$E[\text{cost}] = P(\text{detection}) \times P(\text{conviction}) \times \text{penalty}$$

Because the three terms multiply, a severe penalty applied to almost nobody produces a small expected cost (Figure 6.7b). Raising the probability of detection therefore usually buys more deterrence per unit of effort than raising the sentence, which is the opposite of the political instinct in most jurisdictions. The case funnel makes the point concretely (Figure 6.8a): of the offences committed, only a fraction is detected, a smaller fraction is investigated, a smaller fraction still reaches prosecution; fewer end in conviction. Each narrowing multiplies into the expected cost.

Two further points deserve emphasis. Enforcement without livelihood alternatives displaces rather than removes the activity, since the people at the source of the chain are usually the poorest participants in it and the least able to stop. And enforcement that treats rangers as expendable fails on its own terms: ranger welfare, training, equipment and legal protection are enforcement investments, not welfare extras. Predictive tools that rank locations by risk are useful only when their training data reflect where offences occurred rather than where patrols happened to go, a distinction that is easy to state and hard to honour.

The framing of wildlife crime has changed over the past decade. High-value trafficking is now understood as serious organised crime, which brings a different toolkit into play: financial investigation, asset recovery, anti-money-laundering reporting obligations and the offences of conspiracy and criminal association, all of which target the organisers rather than the couriers who are usually the only people arrested. Corruption is the enabling condition for most large-scale trafficking, since permits, inspections and prosecutions are the points at which a payment changes the outcome, so anti-corruption measures are conservation measures.

Table 6.6. Enforcement interventions along the trade chain.

Stage	Intervention	Evidence of effect	Failure mode
Source	Patrol coverage informed by structured planning	Reduced offtake where coverage is real	Patrols recorded rather than performed
Source	Community information networks	Strong where trust exists	Collapses with a single betrayal of confidence
Transit	Container profiling and detection dogs	Substantial seizure rates at focused ports	Displacement to unmonitored routes
Transit	DNA and isotope forensics	Links seizures to source populations	Requires reference databases
Market	Online marketplace monitoring	Effective where platforms cooperate	Migration to closed channels
Market	Financial investigation of proceeds	Reaches organisers rather than couriers	Requires specialist capacity
Demand	Behaviour change and social norms	Measurable but slow	Evaluated on awareness rather than behaviour




Figure 6.7. Enforcement. (a) The trade chain from source to demand, with the interventions available at each stage. (b) Because expected cost is the product of detection, conviction and penalty, raising detection usually buys more deterrence than raising the sentence.




Figure 6.8. Compliance in practice. (a) The case funnel from offences committed to convictions obtained. (b) Listings accumulate steadily, with marine taxa driving recent growth. (c) The links in the compliance chain that are typically weakest, which is where investment yields most.

6.3 Synthetic Biology and Conservation

6.3.1 A ladder of intervention

Public discussion of this field collapses very different techniques into a single argument. It is more productive to arrange them on a ladder of increasing novelty, uncertainty and irreversibility (Figure 6.9a). At the base sits biobanking, the cryogenic preservation of gametes, embryos and tissue, which alters no organism and which was treated in Section 4.3.2 as an adjunct to pedigree management. Above it sits cloning from preserved cells, demonstrated in a small number of species including the black-footed ferret and the Przewalski's horse, which produces an individual carrying an existing genome rather than a new one. Above that sits genetic rescue by translocation, which is ordinary conservation practice using no engineering at all.

The upper rungs are different in kind. Gene editing of a specific trait, such as engineering disease resistance into a threatened species, remains at laboratory and captive stages with no authorised wild release. Engineered gene drives, which force an allele through a population faster than inheritance would allow, exist only in laboratories; as of the 2025 reviews no field release has occurred anywhere. At the top sits proxy creation, marketed as de-extinction, which produces a modified living species carrying selected traits of an extinct one.

The governance position shifted in 2025. At its World Conservation Congress in Abu Dhabi in October, IUCN adopted its first Policy on Synthetic Biology, carried with about eighty-eight per cent support, while a competing motion for a moratorium on the release of genetically modified wild species was defeated by a narrow margin. The adopted policy requires case-by-case assessment against conservation objectives rather than a blanket rule in either direction. Supporters describe this as evidence-based decision-making; opponents describe the outcome as a victory for commercial biotechnology interests over precaution. Both characterisations are worth taking seriously; a student of the field should read the motion texts rather than the press coverage of either side.

Table 6.7. Applications of synthetic biology in conservation, with their status.

Application	What it would do	Status	Principal concern
Biobanking	Preserves genetic material indefinitely	Routine	Funding continuity; sample provenance
Cloning from cell lines	Returns lost founder genomes to a pedigree	Demonstrated	Cost; low success rates
Disease-resistance editing	Confers resistance, as proposed for chestnut blight and amphibian disease	Laboratory and captive	Ecological effects of the modified trait
Adaptive trait editing	Raises thermal tolerance, as explored for corals	Experimental	Reduces variation while adding one trait
Suppression gene drive	Crashes an invasive population	Laboratory only	Spread beyond the target range; irreversibility
Modification gene drive	Alters a trait without removing the population	Laboratory only	Resistance alleles; containment
Engineered biocontrol	Replaces a chemical control agent	Early	Non-target effects
Proxy creation	Produces an animal resembling an extinct species	Announced, contested	Diverts attention and funds; welfare

Two further considerations frame the debate. The first is dual use. The techniques that could suppress an invasive rodent population could suppress a native one; the capability is not confined to the institutions that developed it, which places synthetic biology alongside other technologies where biosafety and biosecurity oversight is a condition of legitimacy rather than an obstacle to it. The second is how the precautionary principle should be read. Its strong form treats uncertain novel risk as a reason not to act, which protects against unknown harms while accepting the known harms of continued decline. Its weaker form requires proportionate caution, evidence proportionate to the stakes; and reversibility where reversibility is achievable. Most working scientists and most regulators have converged on the weaker reading, which is what the IUCN policy reflects, though the disagreement is about values as much as evidence and will not be settled by more data alone.

6.3.2 Gene drives

A gene drive is a genetic element that is inherited by substantially more than half of an individual's offspring. Because inheritance is biased, an allele can spread through a population even when it lowers fitness, which is the opposite of ordinary selection. The consequence is visible in the trajectory (Figure 6.9b): an allele introduced at one per cent can approach fixation within twenty generations, whereas a neutral allele stays where it was put.

Two designs serve different purposes. A suppression drive spreads an element that reduces fertility or biases the sex ratio, with the aim of crashing the population; the application of interest to conservation is the eradication of invasive rodents from islands, where current methods rely on toxicants. A modification drive spreads a trait through a population without

reducing it, the leading example being mosquitoes rendered unable to transmit malaria parasites. Neither has been released.

The technical objections are as serious as the ethical ones. Resistance alleles arise readily when repair of the cut site produces a sequence the drive can no longer recognise; and because such alleles are strongly favoured they spread rapidly, which can defeat a drive within a few generations. Containment is the deeper problem: a drive designed for an island population will spread to the mainland if a single fertile individual crosses; for a species such as the house mouse that is an unremarkable event. Self-limiting architectures address this directly by building in an exhaustion mechanism, so that drive activity declines after a set number of generations and the element is lost. The trade-off is that a self-limiting drive requires larger and repeated releases, which raises cost while reducing risk.

Governance has not kept pace with the science, though it is moving. Any release would engage the Convention on Biological Diversity and its Cartagena Protocol, national biosafety regulation; and the consent of the communities and states that would host it. The consensus position, reflected in the IUCN policy, is that decisions must be case by case, that the counterfactual matters, since the alternative to a drive is rarely no intervention but rather continued use of toxicants or continued extinction; and that the affected public must be engaged before the technology is ready rather than after.

6.3.3 De-extinction and what it delivers

No technique available now or foreseeable recreates an extinct species. What is achievable is a proxy: a living animal of an extant species, edited to carry a small number of traits inferred from an extinct relative's genome (Figure 6.10a). The distinction is not pedantry, because the conservation claims made for the technology depend on which of the two is delivered.

The case that clarified this for the public was the announcement in April 2025 of three wolf pups presented as de-extincted dire wolves. What was actually done was twenty genetic changes across fourteen genes in a grey wolf genome, followed by cloning and surrogacy. Specialists in canid evolution responded that dire wolves and grey wolves are separated by millions of years of divergence and that hundreds or thousands of changes would be required before the description could be considered, so the animals are best described as modified grey wolves. The episode is instructive on three counts: on the gap between a technical achievement and its marketing, on the willingness of media coverage to relay a claim without checking it; and on the responsibility of scientists to describe their own work accurately.

Three arguments against pursuing proxy creation as a conservation priority are worth stating plainly. The first is opportunity cost (Figure 6.10b): the same expenditure spent on habitat protection or invasive species removal secures far more species-years; conservation budgets are not expandable at will. The second is moral hazard, the concern that a public belief in reversibility weakens the case for preventing extinction in the first place. The third is ecological: a proxy returns without the learned behaviour transmitted by its parents, without its microbiome, into an ecosystem that has reorganised in its absence and that may no longer contain what it needs.

The counterargument deserves a fair hearing. The techniques developed for proxy creation are the same techniques that serve living species: cloning to return lost founder genomes to a pedigree, editing for disease resistance in a species being driven to extinction by a pathogen, biobanking that preserves options. Investment in the flashy application has funded tools with legitimate uses. The reasonable position is neither prohibition nor enthusiasm but insistence that each proposal state what it will deliver, at what cost, against which alternative use of the same money, with which reversibility; and with whose consent.

Table 6.8. Questions a proposed genetic intervention should answer before it is authorised.

Question	Why it matters	A weak answer looks like
What is the conservation objective?	The tool must serve a goal defined without reference to the tool	The objective is to demonstrate the technology
What is the counterfactual?	The comparison is with the best alternative use of the resources	The alternative is described as doing nothing
How far can it spread?	Containment determines whether a decision is reversible	Spread is assumed to remain within the release area
How would it be stopped?	A reversal mechanism must exist before release	Reversal is described as a future research topic
Who consents?	Host communities and states carry the risk	Consultation occurred after the design was fixed
What monitoring follows?	Effects must be detectable early enough to act	Monitoring is funded for one year
Who is liable?	Transboundary effects need an accountable party	Liability is unaddressed




Figure 6.9. Genetic intervention. (a) A ladder from routine preservation to proxy creation, ordered by novelty and reversibility, with the 2025 IUCN policy decision noted. (b) A gene drive biases inheritance so that an allele can approach fixation in a few tens of generations; self-limiting designs are built to fade.

6.4 Synthesising Coexistence: A 2050 Vision

6.4.1 Integrated landscape planning

Every previous chapter contributes a layer to the same landscape (Figure 6.11a). Governance and rights determine who decides. Protected cores and other effective area-based measures secure the places that cannot tolerate use. Corridors and stepping stones keep the cores connected. Working land is managed so that wildlife can pass through and, where possible, live in it. Restoration rebuilds what was lost, with refugia identified and protected against climate

change. Monitoring and finance run underneath everything, because a plan without measurement is an aspiration and a plan without money is a document.

Integration means that the layers are designed together rather than in sequence by different agencies. In practice this requires three things that are institutionally difficult. It requires a single spatial plan that all sectors are bound by, including agriculture, transport and energy, since a corridor is worth nothing if the highway authority is not party to the plan. It requires the plan to be legible on the ground, so that a farmer can say which zone a field is in. And it requires a monitoring system reporting against the plan rather than against each agency's own indicators, which is the only way to detect that the layers have drifted apart.




Figure 6.10. De-extinction assessed. (a) A scale from an unmodified extant species to the extinct organism itself, with the 2025 wolf pups placed near the left-hand end. (b) The opportunity cost argument: indicative species-years secured per unit of expenditure across conservation options.

The difference such a plan makes is easiest to see visually (Figure 6.13). Today's common condition is a set of isolated cores in a hostile matrix, severed by linear infrastructure. The 2050 target condition retains the same cores, links them with corridors and stepping stones, adds crossings where infrastructure intersects movement; and manages the intervening land so that it is permeable rather than empty. Nothing in that picture requires a technology that does not exist. It requires coordination, tenure security and sustained finance, which are harder.

Table 6.9. Layers of an integrated landscape plan, with the chapter that supplies the method.

Layer	What it delivers	Method from	Failure if omitted
Governance and rights	Clear authority and secure tenure	Chapters 4 and 5	Nobody is accountable; rules are ignored
Protected cores and OECMs	Places secure from incompatible use	Chapter 4	No refuge for sensitive species
Connectivity	Movement, gene flow, range shift	Chapter 3	Isolated populations decline to extinction
Working land and coexistence	A permeable matrix people can live in	Chapters 3 and 5	Conflict escalates; the matrix becomes hostile
Restoration and refugia	Recovered habitat and climate havens	Chapter 3	Losses are never reversed
Monitoring and finance	Evidence of change and the means to act	Chapters 2 and 5	Nothing is detected; nothing is sustained




Figure 6.11. The integrated landscape. (a) Six layers that must be designed together rather than sequentially by separate agencies. (b) Four trajectories to 2050: continued decline, intensified pressure, pressures halted; and the recovery implied by bending the curve.

Finance is the constraint that binds most tightly. Estimates of the annual shortfall between what is spent on biodiversity and what would be needed to meet the current global targets run to hundreds of billions of United States dollars; a substantial share of existing public expenditure actively subsidises the drivers of loss through support for agriculture, fisheries and fossil fuels. Redirecting harmful subsidies is therefore the largest single lever available; it is a political rather than a technical problem. New instruments are emerging, including biodiversity credits intended to work as carbon credits do. They deserve the same scepticism that the payment schemes of Chapter 5 required: a credit is worth what its measurement, its additionality and its permanence are worth; a market built on weak versions of all three will transfer money without transferring outcomes.

6.4.2 Climate resilience

Climate change removes the assumption on which most conservation planning rested, namely that the past distribution of conditions predicts the future one. The RAD framework introduced in Section 3.4 organises the response into three honest options (Figure 6.12a). Resist means holding a system in its current state, which is defensible where the value is high and the effort is affordable, though it becomes progressively more expensive as conditions diverge. Accept means allowing change to proceed without intervention, which is the right answer far more often than agencies find comfortable to say. Direct means steering change toward a chosen alternative state, which requires being explicit about what that state is.

Choosing among the three requires a judgement about how far conditions will move, which is uncertain (Figure 6.12b). Retention of suitable range falls steeply with warming; the species that lose most are those with narrow ranges and no upslope or poleward escape. That uncertainty argues for actions that perform well across scenarios rather than actions optimised for one (Figure 6.12c). Protecting climate refugia, restoring connectivity, managing for heterogeneity so that microclimates persist within a site; and reducing non-climate stressors all pay off whichever trajectory materialises. Assisted gene flow and assisted colonisation are more powerful and more scenario-dependent, so they belong later in the sequence rather than first.

The framing that matters most for the coming decades is the difference between halting decline and reversing it (Figure 6.11b). Stopping the pressures stabilises biodiversity at a diminished level. Recovery requires restoration at a scale that is currently rhetorical rather than budgeted. The gap between those two curves is the practical content of the phrase bending the curve; closing it is a question of finance and land-use policy rather than of ecological knowledge.




Figure 6.12. Climate adaptation. (a) The RAD framework: resist, accept or direct, ordered by departure from historical conditions. (b) Range retention falls steeply with warming, most severely for restricted species. (c) Actions ranked by how well they perform across climate scenarios rather than under one.




Figure 6.13. Two landscapes. (a) The common condition today: isolated cores in a hostile matrix, severed by infrastructure. (b) The 2050 target: the same cores linked by corridors and stepping stones, with crossings where infrastructure intersects movement and a managed, permeable matrix between them.

Monitoring has to change with the ambition. Reporting against a global framework requires variables that are comparable between countries, which is the purpose of the essential biodiversity variables: a small set of measurements spanning genetic composition, species populations, species traits, community composition, ecosystem function and ecosystem structure, chosen so that national systems can feed a global picture without being identical. The practical implication for a national agency is that a monitoring programme should be designed to produce those variables from the outset rather than to produce a report; and that the sensor networks of Section 6.1 should be treated as the means of producing them at a cost a country can sustain. A monitoring system that cannot detect a change of the size that would trigger a management response is not a monitoring system, whatever it costs.

6.4.3 Working in this field

The competencies this book has drawn on are unusually broad. A single conflict programme required population ecology, sampling design, engineering, economics, law and facilitation. A single protected area required spatial optimisation, governance analysis, budgeting and negotiation. No individual holds all of it, which is why the useful shape for a career is the one often described as T-shaped (Figure 6.14a): one discipline held to genuine depth, supported by enough breadth in the others to work with the people who hold them.

Certain skills are transferable across every path in the field. Field craft and the judgement to work safely remain foundational. The ability to design a study before collecting data prevents most of the failures described in Chapter 2. Coding and reproducible analysis have moved from optional to expected. Writing and speaking clearly determine whether work influences

anything. And the capacity to work respectfully across cultures and knowledge systems, treated in Section 5.3, is the difference between a plan that is implemented and one that is resented.

The paths themselves are more varied than the stereotype of the field biologist suggests (Figure 6.14b, Table 6.10). Research and academia remain one route. Management agencies employ the largest number of practitioners in most countries. Non-governmental organisations span field implementation, technical advice and advocacy. Policy and law offer leverage that field work rarely matches. Technology and data roles have grown quickly and are undersupplied. Community-facing practice, including Indigenous ranger programmes, is where a great deal of the most effective work now happens. Movement between these paths is common and increasingly expected.

Table 6.10. Career paths in wildlife biology and conservation, with what each demands.

Path	Typical roles	Distinctive skill	Main frustration
Research and academia	Field ecologist, quantitative ecologist, modeller	Study design and inference	Funding cycles; slow translation into practice
Management agency	Protected area manager, wildlife officer, planner	Operational judgement under constraint	Political direction; limited budgets
Non-governmental sector	Programme manager, technical adviser, campaigner	Delivery within a funded project	Short project horizons
Policy and law	Negotiator, regulator, environmental lawyer	Drafting and negotiation	Distance from the field
Technology and data	Conservation technologist, remote sensing analyst	Building tools others can use	Solutions in search of problems
Community practice	Extension officer, facilitator, Indigenous ranger	Trust built over years	Recognition and pay lag the contribution




Figure 6.14. Working in conservation. (a) The T-shaped professional: depth in one discipline with enough breadth to collaborate across the others. (b) The principal career paths, which practitioners increasingly move between rather than choosing once.

A closing observation. The tools described in this chapter are powerful; it would be easy to read the sequence of chapters as a march from ignorance toward technical mastery. That reading

would be wrong. The constraints on conservation outcomes are rarely technical. We can count animals we cannot protect, map habitats we cannot secure and sequence genomes of species we are losing. What determines the outcome is whether the people who live with wildlife have reason to keep it, whether institutions do what they promised; and whether the costs of coexistence are shared by those who benefit from it. Technology serves those ends when it is embedded in the pipeline of Section 6.1 all the way through to a decision. It serves nothing when it stops at a dashboard.

Chapter Summary

- Conservation technology is a pipeline from sensor to decision; automating detection while neglecting estimation buys speed without knowledge.
- Data volume has outgrown human review capacity, which is what made automated classification necessary rather than merely attractive.
- Detection, species classification and individual identification are distinct tasks with sharply different performance; a stated accuracy is meaningless without naming the task.
- Precision and recall trade off against one another; the operating threshold is a management decision rather than a technical one.
- Domain shift silently degrades a model deployed in a new region; validation must hold out whole sites; and a human verification sample must be permanent.
- A model output is a count, not an abundance estimate; detection probability must still be estimated.
- No remote sensing platform is best; each occupies a position in the trade-off between resolution, extent and revisit interval.
- Tag mass limits battery capacity, so temporal resolution is paid for directly in tag life; accuracy across tracking systems spans three orders of magnitude.
- Conservation law is layered from global treaties to the field; the binding problem is implementation rather than ambition.
- CITES regulates international trade through three appendices, with the non-detriment finding as its scientific core; CoP20 in Samarkand in 2025 newly listed or uplisted nearly ninety animal species.
- Expected cost of offending is the product of detection, conviction and penalty, so detection capacity usually buys more deterrence than severity.
- Genetic interventions form a ladder from biobanking to proxy creation; IUCN adopted its first Policy on Synthetic Biology in October 2025 and rejected a moratorium.
- Gene drives bias inheritance so an allele can approach fixation in tens of generations; none has been released; resistance alleles and containment remain unsolved.
- De-extinction delivers a proxy rather than a species; the 2025 wolf pups carried twenty changes across fourteen genes in a grey wolf genome.
- An integrated landscape plan aligns governance, cores, connectivity, working land, restoration and finance; bending the curve to 2050 is a question of policy and money rather than of ecological knowledge.

Key Terms

Domain shift the loss of model accuracy when data differ systematically from the training set

Duty cycling scheduling a tag to record or transmit intermittently in order to extend its life

F1 score the harmonic mean of precision and recall

Gene drive a genetic element inherited by substantially more than half of an individual's offspring

Ground sample distance the ground area represented by one pixel of an image

Implementation gap the distance between an agreed commitment and what happens on the ground

Non-detriment finding the scientific judgement that an export will not harm the survival of the species

OECD an other effective area-based conservation measure delivering biodiversity outcomes outside formal protection

Precision the proportion of predicted detections that are correct

Proxy species a living organism engineered to carry selected traits of an extinct one

RAD framework the choice among resisting, accepting or directing ecological change

Recall the proportion of true events that are detected

Resistance allele a sequence variant that a gene drive can no longer recognise or cut

Self-limiting drive a drive architecture designed to lose activity after a set number of generations

T-shaped professional depth in one discipline combined with working breadth across others

Review Questions

1. A classifier reports 97 per cent accuracy on a camera-trap dataset in which 95 per cent of images are empty. Explain why this figure is uninformative, then state which measures you would ask for instead.
2. A model detects 1,840 individuals of a species across a camera array. Explain, with reference to Chapter 2, why this number should not be reported as a population estimate.
3. You are deploying a model trained in southern Africa to a site in South Asia. Set out the validation and retuning steps you would require before any output is used in a decision.
4. A team proposes counting nesting seabirds on a 40-hectare island. Compare a drone survey with a commercial satellite tasking on resolution, cost, disturbance and repeatability.
5. A 180 g bird is to be tracked through two migrations. Apply the tag mass rule, then explain the consequences of your mass limit for fix frequency and data return.
6. A range state proposes exporting a listed reptile under Appendix II. Set out the evidence its scientific authority would need before a non-detriment finding could be made.
7. A government doubles the maximum sentence for wildlife trafficking while its detection rate remains at four per cent. Estimate the change in expected cost to an offender, then propose an alternative allocation of the same budget.

8. A proposal is made to release a suppression gene drive against invasive rodents on an inhabited island. Using Table 6.8, identify the three questions you would expect to be hardest to answer; explain why.

Further Reading

Tuia, D. and colleagues. Perspectives in machine learning for wildlife conservation, in Nature Communications, for the framing used in Section 6.1.

CITES Secretariat. The text of the Convention together with the CoP20 outcomes from Samarkand, 2025.

Secretariat of the Convention on Biological Diversity. The Kunming-Montreal Global Biodiversity Framework and its monitoring framework.

IUCN. Policy on Synthetic Biology in Relation to Nature Conservation, adopted at the World Conservation Congress in Abu Dhabi in October 2025.

Esvelt, K. M. and colleagues on gene drive architecture and containment, together with the National Academies report on gene drives in the environment.

Watson, J. E. M. and colleagues on protected area effectiveness and integrated landscape planning, for the synthesis in Section 6.4.

About Editors



Dr. Ranjit R. Raut is currently working as a Young Scientist and Biodiversity Researcher at Department of Zoology, Yogeshwari Mahavidyalaya Ambajogai, affiliated to Dr. Babasaheb Ambedkar Marathwada University Chhatrapati Sambhajnagar, Maharashtra, India. Dr. Ranjit has completed Diploma in Wildlife Biology from Kongunadu College (Autonomous) Coimbatore, Tamil Nadu affiliated to Bharathiar University Coimbatore. He has almost 12 years of research experience. He has published 63 research papers in reputed, leading National and International Journals. He has registered 7 Indian Patents. He is a life member of Biotech Research Society India. He is a board member of Fireflies Asian Association. He has worked as a volunteer at Wildlife Crime Control Bureau (WCCB, MoEFCC). He has received many International, National and State level awards like Young Scientist, Best Researcher, Excellence in Environmental Sustainability, Friend of Animal etc. He has presented and participated in many National and International workshops, seminars, conferences and symposia. His areas of research interest are Wildlife Biology, Biotechnology, Life Sciences and Biodiversity Conservation.



Dr. Komal L. Jagzap is a Zoologist, Researcher and Academician with a strong academic and research background in Life Sciences. She has been actively involved in teaching postgraduate and undergraduate students, guiding academic research, and participating in scientific and educational initiatives. Her research work focuses on understanding biological diversity, insect biology, pest management and conservation-oriented approaches. Through her teaching and research activities, she aims to bridge fundamental biological concepts with contemporary research methodologies and their practical applications. She is passionate about developing scientific thinking, research aptitude and a deeper understanding of biological sciences among students and young researchers.



Dr. Yogesh V. Bidwe is an enthusiastic academician and researcher. Currently working as a Researcher at Department of Zoology, Dr. Rafiq Zakaria College for Women Chhatrapati Sambhajnagar. affiliated to Dr. Babasaheb Ambedkar Marathwada University Chhatrapati Shambhajnagar. He has 12 years working experience in the diverse fields of Life Sciences. He has completed his Ph.D and Post-Graduation at Department of Zoology, Dr. Babasaheb Ambedkar Marathwada University Chhatrapati Sambhajnagar. He has published 03 research articles in Web of Science Indexed peer reviewed National & International journals.



Kumari Shilpa Patel is a fish and fisheries science researcher from Barakhurd, Bihar, currently based in Patna. She holds a B.Sc. and M.Sc. in Zoology from Magadh University and is pursuing a Ph.D. at Gaya College. Kumari Shilpa Patel has published 12 research papers in reputed journals, including three in Scopus-indexed journals and two in Web of Science and others in reputed peer review journals. She has contributed to five book chapters and presented her research at various national and international conferences, including the Indian Science Congress in 2023. Her goal is to leverage her knowledge and research skills for organizational growth.

MRP: Rs. 299/-

ISBN: 978-81-18-05442-9



9 788118 054429