The Economics of Contracting for Conservation

WORKING PAPER: DO NOT CITE OR CIRCULATE WITHOUT PERMISSION

Bryan Leonard Kailin Kroetz

Abstract

Private-land conservation has moved steadily from fee-simple acquisition toward conservation easements and, more recently, toward short-term payments for ecosystem services (PES). To date, economics offers no unified account of when a conservationist should buy land, buy development rights, or engage in PES. We provide one, modeling a conservationist who minimizes the cost of meeting a habitat target on working lands whose profit-maximizing owners face stochastic development pressure. On a single parcel, PES weakly dominates acquisition, and the easement decision turns on the value of preventing development rather than its probability. Across a landscape, the convexity of incentive costs leads the conservationist to spread PES thinly over many parcels, and easements take on a second role: preserving parcels over which to spread that effort. Development pressure therefore raises both the cost and the value of easements, and its effect on easement use can be non-monotonic. Per-contract monitoring costs restore acquisition and pure easement programs as optimal instruments, and their decline—driven by remote sensing and cheap verification—offers an account of the century-long drift from owning land to renting habitat.

1 Introduction

Land conservation is central to many of the most pressing global environmental challenges, including biodiversity loss, climate change, food security, and energy production (Meyfroidt et al., 2022). Human land uses have altered more than 80% of the Earth's terrestrial surface, and much of the remainder is under threat from development, deforestation, and agricultural intensification (Lambin and Meyfroidt, 2011; Meyfroidt et al., 2022). This concern has prompted calls to expand protected areas and deepen existing protections to secure 30% of the Earth's land and waters by 2030 under the "30 by 30" initiative, with proposals to reach 50% by 2050 (Baillie and Zhang, 2018). That agenda is rooted in a "fortress" model of conservation that emphasizes public ownership and formal protection, through mechanisms such as the designation of Wilderness Areas and the public acquisition of private land.

Conservation on private land, by contrast, has moved steadily away from fee-simple acquisition and toward contractual arrangements that stop short of a full transfer of ownership (Parker, 2004). These include conservation easements, which pay landowners to extinguish development rights, and a more recent class of short-term arrangements—"habitat leasing" and other incentive payments that fall under the heading of payments for ecosystem services (PES) (Salzman et al., 2018). The shift reflects a growing recognition that private working lands—farms, ranches, and private forests—hold substantial conservation value, providing habitat, carbon storage, and stopover sites and corridors for migratory species that do not remain within traditional protected areas (Kremen and Merenlender, 2018; Byrd et al., 2019). Often a small change in land use delivers a large conservation gain. For example, the Nature Conservancy's BirdReturns program pays Central Valley farmers to flood fields at the moments migratory birds need them, rather than seeking permanent land use change.

This progression from acquisition to easements to PES raises a question that the economics literature has not addressed directly: facing a landowner whose profit-maximizing choices determine the conservation outcome, which instrument should a conservationist use, and when? We develop a model to answer it. A conservationist seeks to secure conservation on private working land at least cost, choosing among fee-simple acquisition, conservation easements, and PES, and combining them across a landscape subject to a habitat target. The model combines a profit-maximizing landowner and a cost-minimizing conservationist facing stochastic development to place the instruments in a single framework and yields a sequence of results that speak to the observed shifts in conservation practice. We build the analysis in three steps. First, we characterize the conservation cost function on a single parcel. Then, we consider the optimal portfolio of strategies on a landscape containing N parcels. Finally, we introduce landowner heterogeneity and transaction costs to the baseline model.

On a single parcel, PES weakly dominates fee-simple acquisition, because a landowner paid to adjust inputs continues to earn profit from the parcel and so secures any conservation outcome at no greater cost than buying the land. Similarly, a pure easement strategy is strictly dominated on a single parcel because it is always optimal to do a small amount of PES. The decision to add an easement to a PES contract depends on the conservation value of preventing development weighed against the landowner's gain from developing, and not on the probability that development occurs, which scales the easement's cost and its benefit equally and so cancels from the comparison. These single-parcel results sharpen and generalize the intuition in the prior literature comparing acquisition to easements (Parker, 2004; Shah and Ando, 2016).

The landscape model is where the central mechanism emerges. Because the cost of incentive payments is convex in the size of the adjustment asked of a landowner, the conservationist minimizes cost by spreading PES thinly across many parcels rather than concentrating it on a few. This spreading logic reduces a multi-parcel, multi-instrument problem to a single choice—how many easements to buy—and reveals a function of easements that the single-parcel analysis misses, because an easement preserves a parcel against development and that parcel is then available to absorb a share of the spread-out PES effort. Easements therefore provide two benefits at the landscape scale, preventing development and enlarging the pool over which conservation is spread in the future. As a consequence, development pressure has an ambiguous, potentially non-monotonic effect on the optimal number of easements, because the same rise in pressure makes easements both more costly and more valuable. We show this formally and characterize when each force dominates, and we show that the optimal policy is distribution-free in a precise sense—it depends on the development-offer distribution only through its conditional mean.

Two extensions connect the model to the patterns observed across real landscapes. Allowing landowners to differ generates a targeting logic where the conservationist directs PES toward parcels that are cheap to adjust and easements toward parcels that are cheap to protect or expensive to reach through PES, a sorting that matches the concentration of easements on development-prone ranchland and of PES on compatible cropland. Introducing a per-contract monitoring cost rehabilitates the two instruments that the frictionless model rules out. When monitoring is expensive, fee-simple acquisition can become the low-cost way to secure large input reductions, and pure easement programs—of the kind run by the USDA and many states—can become optimal. Because the long decline in monitoring costs, driven by remote sensing and crowd-sourced data, shifts the model's regime boundaries, it offers an account of the historical drift from acquisition toward easements and then toward PES.

Our contribution is to provide the first unified analysis of fee-simple acquisition, conservation easements, and PES as alternative instruments for the same conservation objective. Existing work treats these mechanisms largely in isolation: there is a large literature on the design of PES programs (Jack et al., 2008), an established literature on the optimal selection of land for acquisition across a landscape (Conrad et al., 2012), and analyses comparing acquisition to easements (Parker, 2004) or permanent to temporary protection under development threat (Shah and Ando, 2016), but none that places all three instruments in a common framework. In doing so we connect conservation to the theory of the firm. Just as economics long lacked a theory of why some transactions occur in markets and others within firms until Coase (1937) and later Hart (1988); Hart and Moore (1990), conservation has lacked a comparable account of why organizations choose among ownership, partial property rights, and contracts. This paper is a step toward such a "theory of the conservationist." The remainder proceeds as follows. Section 2 sets up the model; Sections 3 and 4 present the single-parcel and landscape results; Section 5 develops the heterogeneity and monitoring extensions; and Sections 6 and 7 discuss implications and conclude.

2 Model Setup

We model conservation as a contracting problem between a single landowner, L, and a conservationist, C, who wishes to secure additional conservation on the landowner's working land. We first describe the landowner's profit-maximization problem, which pins down the baseline level of conservation provided in the absence of any contract, as well as the landowner's opportunity cost of engaging in conservation. We then describe the conservationist's objective, which we cast as the minimization of the cost of achieving a conservation target. Finally, we introduce development pressure and define the cost of each conservation instrument. Throughout, we assume that L holds the property right to the land and is free to choose all inputs to maximize profit, so that any change in land use desired by C must be paid for. We focus on the minimum payment that would make L just indifferent to participating—that is, we assume L has no bargaining power and absorbs no surplus.[1]

2.1 The Landowner

Consider a single parcel within a working landscape—for example, a farm or ranch in a region dominated by agriculture—owned by a landowner whose objective is to maximize profit. Prior to any conservation contract, L chooses two inputs, x_1 and x_2, to solve

$$\max_{x_1,x_2} \quad \pi(x_1,x_2) = p\,f(x_1,x_2) - w_1x_1 - w_2x_2, \tag{1}$$

where p is the output price, w_i is the price of input i, and $f(x_1,x_2)$ is a production function that we assume to be non-negative, increasing, and concave in both inputs. We impose $x_i \geq 0$. The landowner's optimum is characterized by the first-order conditions

$$p\,f_1(x_1,x_2) - w_1 = 0, \qquad p\,f_2(x_1,x_2) - w_2 = 0, \tag{2}$$

with solution (x_1^L, x_2^L). We denote the landowner's maximized profit by $\pi^L \equiv \pi(x_1^L, x_2^L)$.

We assume that L's choice of x_1 has no consequence for conservation, while L's choice of x_2 completely determines the conservation outcome. For instance, x_2 might be the stocking density of cattle on a ranch or the acreage of native forage left for pollinators on a farm, while x_1 is a conservation-neutral input such as labor. More generally, x_2 could be a vector of conservation-relevant inputs and x_1 a vector of conservation-neutral inputs; we work with scalars to simplify the exposition. We assume that if C acquired the land and wished to adjust x_2, she would face the same marginal input cost w_2 as the landowner, so that setting $x_2 = 0$ would entail no input cost.[2]

2.2 The Conservationist

The conservationist derives a benefit from the conservation-relevant input on the parcel, summarized by a benefit function $B(x_2)$ that we assume to be concave. We allow B to be increasing or decreasing, depending on the interpretation of x_2: it may be monotonically decreasing (e.g., for pesticide application), monotonically increasing (e.g., for pollinator forage), or single-peaked (e.g.,

[1]This is the appropriate assumption when C faces many potential parcels and none is pivotal, so that competition among landowners drives the payment down to willingness-to-accept. Allowing L to capture a constant share of the surplus would rescale payments but would not alter the contract comparisons below, provided L's bargaining power does not differ across instruments.

[2]Adjustment costs—for example, the cost of reducing x_2 from the level chosen by L—would be borne by one party or the other under any contract and so do not affect the comparisons below.

for tree density in a grassland).[3] We normalize $B(0) = 0$, so that a parcel converted to a developed use—which produces $x_2 = 0$—contributes nothing to C's objective. This normalization links the single-parcel and landscape models below: developed parcels simply drop out of the conservation constraint.[4]

We assume that C allocates a fixed budget across the focal project and many other conservation projects, so that the organization's budget is large relative to the cost of any single parcel but a dollar spent here carries an opportunity cost of a dollar spent elsewhere. This is the situation of a large conservation organization or agency—the Nature Conservancy, say, or the U.S. Department of Agriculture—for which any one parcel is a small part of a much larger portfolio. Under these conditions, the conservationist's problem is equivalent to minimizing the cost of achieving a given conservation target. Letting $\kappa(x_2)$ denote the cost to C of securing x_2 units of the conservation input on the parcel, and $\bar{B}$ a target level of conservation benefit, C solves

$$\min_{x_2} \quad \kappa(x_2) \qquad \text{s.t.} \qquad B(x_2) \geq \bar{B}. \tag{3}$$

The multiplier on the constraint from the associated Lagrangian, λ, is the marginal value to C of an additional unit of conservation benefit—the amount C would pay, at the margin, to relax the target. We use the cost-minimization formulation throughout because it extends to the landscape problem of Section 4, where C must reach a landscape-wide habitat target across many parcels. On a single parcel, problem (3) is the dual of the benefit-cost problem $\max_{x_2} B(x_2) - \kappa(x_2)$, with λ equal to the marginal benefit $B'(x_2)$ at the optimum; we move between the two formulations as convenient, valuing conservation on a parcel at its shadow value λ.[5]

What distinguishes the instruments is the shape of the cost function $\kappa(x_2)$, which depends entirely on the institution through which C engages in conservation. We characterize $\kappa(x_2)$ under each available contract below.

[3] Concavity rules out a convex ("U-shaped") benefit function.

[4] Allowing $B(0) < 0$ would capture additional damages from development, such as pollution or habitat fragmentation, and would only strengthen the case for instruments that prevent development. We abstract from this channel for parsimony.

[5] The dual is the quasi-linear utility formulation $U(x_2) = B(x_2) + y$ subject to $\kappa(x_2) + y = K$, where y is spending on other projects and K is the budget. Because the budget is large, the marginal value of funds is one and the interior optimum satisfies $B'(x_2) = \kappa'(x_2)$. The cost-minimization and benefit-cost formulations therefore deliver identical contract comparisons.

2.3 Development Pressure

A defining feature of working-land conservation is that the land may be converted to a developed use that forecloses conservation. We capture this threat with a two-period timing. In period 1 the conservationist writes contracts; between periods, each parcel receives a stochastic development offer d drawn from a distribution $F(d)$ with support $[0, \bar{D}]$; in period 2 conservation is delivered on the parcels that remain. The landowner sells to the developer—foreclosing conservation—whenever the offer exceeds the value of continuing in the working use, $d > \pi^L$. Payoffs realized in period 2 are discounted by a factor β. The *probability of development* is therefore

$$\theta = 1 - F(\pi^L), \tag{4}$$

which is endogenous to the profitability of the working use: more profitable land is less likely to develop.

A conservation easement extinguishes the development option by paying L for the expected profit forgone by agreeing never to sell to a developer. The *easement cost* per parcel is

$$C_E = \mathbb{E}\big[\max(d - \pi^L,\, 0)\big] = \int_{\pi^L}^{\bar{D}} (d - \pi^L)\, dF(d) = \theta \cdot \mathbb{E}\big[d - \pi^L \mid d > \pi^L\big]\,. \tag{5}$$

The easement cost decomposes into the probability that development occurs times the expected gain from development conditional on its occurring. A key advantage of this formulation is that a single primitive—the upper bound $\bar{D}$ of the offer distribution, which we interpret as the intensity of development pressure—simultaneously governs both the *probability* of development through θ and its *cost* through C_E. This dual role distinguishes our formulation from one in which the probability and value of development are free parameters, and it becomes consequential in the landscape analysis of Section 4.

To obtain closed-form results we will at times specialize to a uniform offer distribution, $d \sim U[0, \bar{D}]$, under which

$$\theta = \frac{\bar{D} - \pi^L}{\bar{D}}, \qquad \mathbb{E}\big[d - \pi^L \mid d > \pi^L\big] = \frac{\bar{D} - \pi^L}{2}, \qquad C_E = \frac{(\bar{D} - \pi^L)^2}{2\bar{D}}. \tag{6}$$

Appendix C shows that the core results do not depend on the uniform assumption: the same objects and comparative statics arise under Beta, log-normal, and Pareto offer distributions.

2.4 Conservation Instruments and Their Cost Functions

The conservationist has four strategies available, each of which gives rise to a different cost function $\kappa(x_2)$ but leaves the underlying problem (3) unchanged.

(i) **No action.** C relies on the landowner to provide x_2^L and spends her entire budget elsewhere. This delivers conservation benefit $B(x_2^L)$ at zero cost to C.

(ii) **Fee-simple acquisition.** C purchases the parcel and then provides her desired level of x_2 at marginal cost w_2. The purchase price is the profit forgone by L, so the cost function is

$$\kappa(x_2) = \pi^L + w_2 x_2, \tag{7}$$

a fixed cost of π^L plus a linear variable cost. Because acquisition also extinguishes the development option, its full cost includes the present value of that option; this term enters acquisition and easement-plus-PES identically and so cancels from the comparison between them.

(iii) **Payment for ecosystem services (PES).** C pays L to adjust x_2 away from x_2^L. The payment is L's willingness to accept, equal to the resulting loss in profit:

$$\gamma(\bar{x}_2) = \pi^L - \pi(\tilde{x}_1(\bar{x}_2),\, \bar{x}_2)\,, \tag{8}$$

where $\tilde{x}_1(\bar{x}_2)$ is L's profit-maximizing choice of x_1 subject to the constraint $x_2 = \bar{x}_2$. The PES cost function is $\kappa(x_2) = \gamma(x_2)$; its shape, which we characterize in Section 3, is what distinguishes incentive payments from the other instruments.

(iv) **Easement plus PES.** An easement is a payment to extinguish the development option, which C may pair with a PES contract to adjust x_2. Its cost is the easement payment C_E plus the PES cost:

$$\kappa(x_2) = C_E + \gamma(x_2). \tag{9}$$

Because the easement guarantees the parcel against development, it is the only instrument other than acquisition that secures conservation in the face of development pressure.

The remainder of the paper characterizes when each of these strategies minimizes the cost of conservation—first on a single parcel (Section 3), then across a landscape with a habitat target (Section 4), and finally under landowner heterogeneity and monitoring costs (Section 5).

3 Single-Parcel Results

We begin with a single parcel, which delivers two results that anchor the rest of the paper: payments for ecosystem services always dominate outright acquisition, and the decision to add an easement turns on the costs and benefits of preventing development rather than on the probability that development occurs. Both results rest on the shape of the PES cost function, which we characterize first.

3.1 The PES Cost Function

The cost of a PES contract is the landowner's willingness to accept the contracted input level, $\gamma(\bar{x}_2) = \pi^L - \pi(\tilde{x}_1(\bar{x}_2), \bar{x}_2)$ from Equation (8). Its shape is what makes incentive payments cheap relative to acquisition.

Lemma 1. *The PES cost function $\gamma(\bar{x}_2)$ is non-negative and strictly convex, attains a unique minimum of zero at $\bar{x}_2 = x_2^L$, and has slope bounded above by w_2: $\gamma'(x_2^L) = 0$ and $\gamma'(\bar{x}_2) \leq w_2$ for all $\bar{x}_2$.* (Proof: Appendix B.)

The logic is an application of the envelope theorem. Because L re-optimizes x_1 for any contracted $\bar{x}_2$, the term in γ that runs through x_1 vanishes at L's constrained optimum, leaving

$$\gamma'(\bar{x}_2) = w_2 - p\,\frac{\partial f(\tilde{x}_1, \bar{x}_2)}{\partial \bar{x}_2}. \tag{10}$$

At $\bar{x}_2 = x_2^L$ this is zero by the landowner's own first-order condition, so the landowner is, to first order, indifferent to a marginal change in x_2: *small adjustments around the working-land optimum are nearly free*. Differentiating once more gives $\gamma''(\bar{x}_2) = -p\,\partial^2 f/\partial \bar{x}_2^2 > 0$ by concavity of f, so γ is convex; and since $\partial f/\partial \bar{x}_2 \geq 0$, the slope never exceeds w_2. The full argument, including the second-order terms, is in Appendix B.

Two features of γ shape the results below. Convexity means it is cheaper to ask many landowners for small adjustments than a few for large ones. The zero slope at x_2^L means the marginal cost of *some* PES is zero, so the first unit of conservation beyond the landowner's optimum is free.

3.2 PES Dominates Fee-Simple Acquisition

The first result is a cost comparison that makes no reference to the benefit function.

Proposition 1. *For any level of the conservation input, the cost of achieving it through PES (possibly combined with an easement) is weakly lower than the cost of fee-simple acquisition. Consequently C never strictly prefers to purchase the parcel.* (Proof: Appendix B.)

The comparison is between $\gamma(x_2)$ and the acquisition cost $\pi^L + w_2 x_2$ from Equation (7). The landowner always retains the option of shutting down conventional production and supplying the contracted x_2 alone, which bounds the PES cost above by $\pi^L + w_2 x_2$; combined with $\gamma(0) \leq \pi^L$ and the slope bound from Lemma 1, this gives $\gamma(x_2) \leq \pi^L + w_2 x_2$ for all $x_2 \geq 0$. PES therefore achieves any given input level at weakly lower cost, and the advantage is generically strict because C may also select a different, cheaper input level under PES. The development option does not disturb the comparison, because its expected cost enters both acquisition and easement-plus-PES identically and so cancels (Appendix B). Figure 1 plots the two cost functions; the PES schedule lies weakly below the acquisition line everywhere, touching it only in the limiting case where production is shut down entirely.

Figure 1. The Cost of Obtaining x_2 Under PES Versus Acquisition

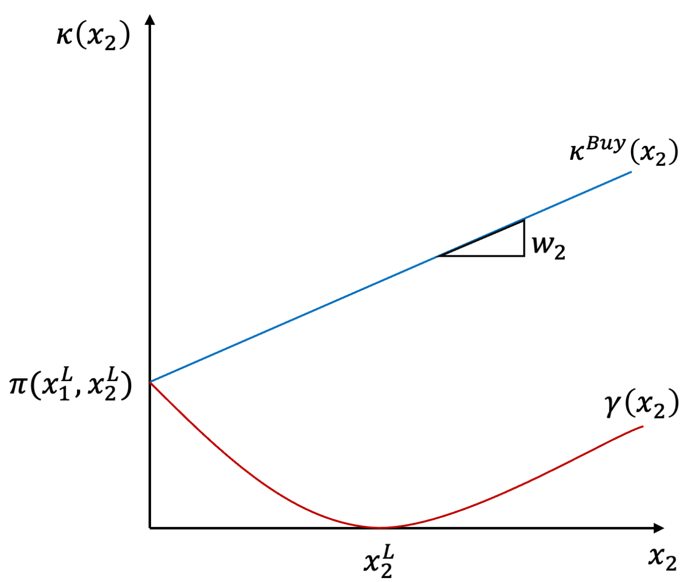

The comparison runs against the emphasis on public acquisition in the "30 by 30" agenda: whenever a landowner will continue to operate the parcel, paying that landowner to adjust inputs is cheaper than buying the land outright. We provide nuance to this finding in Section 5, where we show that once monitoring costs make incentive contracts expensive to administer, acquisition can re-emerge as the low-cost option.

3.3 Easements Versus PES on a Single Parcel

With acquisition ruled out, C's remaining choice on a single parcel is whether to pair PES with an easement that guarantees the parcel against development.[6] The benefit of conservation now re-enters, because an easement is worth purchasing only if the conservation secured by preventing development justifies the cost of extinguishing the development option.

Proposition 2. *C adds an easement to a PES contract on a single parcel if and only if*

$$\beta \underbrace{\left[B(x_2^P) + \pi(\tilde{x}_1, x_2^P) - \pi^L\right]}_{\text{conservation surplus}} > \underbrace{\mathbb{E}\left[d - \pi^L \mid d > \pi^L\right]}_{\text{expected development surplus}}, \tag{11}$$

where x_2^P is the optimal PES input level. The probability of development θ does not enter the condition. (Proof: Appendix B.)

The reason θ drops out is that it scales the easement's cost and its benefit in equal measure. Without an easement, conservation is secured only when development does not occur, an event of probability $1-\theta$, so the expected discounted payoff is $\beta(1-\theta)\left[B(x_2^P)-\gamma(x_2^P)\right]$. With an easement, C pays C_E in period 1 and secures the discounted conservation payoff $\beta\left[B(x_2^P) - \gamma(x_2^P)\right]$ for certain. Subtracting, the easement is preferred when $\beta\theta\left[B(x_2^P) - \gamma(x_2^P)\right] > C_E = \theta\,\mathbb{E}[d - \pi^L \mid d > \pi^L]$, and the common factor θ cancels. Substituting $\gamma(x_2^P) = \pi^L - \pi(\tilde{x}_1, x_2^P)$ yields condition (11).

The interpretation is that C compares two surpluses. The left-hand side is the discounted social surplus from keeping the parcel in conservation use—the habitat benefit plus the agricultural profit earned under the PES contract, net of the baseline profit the landowner would have earned anyway. The right-hand side is the private surplus the landowner would capture from developing, conditional on a development offer arriving. C buys the easement when conservation is worth more than development at the margin, and the likelihood of the development threat materializing affects only the magnitude of the stakes, not their direction. Easements are therefore most attractive on land that is highly valuable for conservation but only modestly valuable for development.

The single-parcel comparison is simple because C has no margin other than this one parcel, and the comparative static on development pressure is correspondingly unambiguous: under the uniform specialization, condition (11) becomes $\beta\left[B(x_2^P) + \pi(\tilde{x}_1, x_2^P) - \pi^L\right] > (\bar{D} - \pi^L)/2$, so a

[6] A *pure* easement—protecting the parcel but leaving land use at the baseline x_2^L—is the $x_2 = x_2^L$ corner of this strategy. Because $\gamma'(x_2^L) = 0$, a marginal amount of PES is costless, so easement-plus-PES weakly dominates the pure easement, strictly whenever $B'(x_2^L) \neq 0$. Pure easements become strictly optimal only once the monitoring costs of Section 5 make even trivial PES expensive to administer.

rise in development pressure $\bar{D}$ raises the conditional development surplus and makes easements unambiguously *less* attractive. As we show next, this result does not survive the move to a landscape, where preserving a parcel through an easement also yields a benefit that the single-parcel problem does not capture: an additional parcel across which to spread PES in the future.

4 Landscape Conservation

Conservation organizations rarely contract over a single parcel in isolation. They face a landscape of many parcels and a habitat target to be met across all of them, and the central question becomes not *whether* to act on a given parcel but *how to allocate* conservation instruments across parcels to achieve a landscape-scale objective at the lowest cost. We now embed the single-parcel problem in a landscape, describing the pool of parcels available for conservation, characterizing the optimal allocation of PES across that pool, reducing the conservationist's problem to a single choice variable—the number of easements—and then developing the comparative statics of that choice and the map of optimal instrument regimes they imply.

4.1 Setup

The landscape consists of N identical parcels, each with the landowner optimum x^L (we drop the subscript on x_2^L for clarity, writing $x^L \equiv x_2^L$). The conservationist must deliver a landscape-wide habitat target $\bar{B}$ and minimizes the total cost of doing so. At the landscape scale we take habitat to be additive across parcels and linear in the conservation input, so that aggregate habitat is $\sum_i x_i$ and the target is expressed in units of that input.[7] The two-period timing of Section 2.3 orders the instruments: in period 1 C chooses how many easements to purchase, and in period 2, after each parcel has drawn its development offer $d \sim F$ on $[0, \bar{D}]$ and developed if $d > \pi^L$, C writes PES contracts on whatever parcels remain.

By the development assumption of Section 2.3, a developed parcel produces $x_2 = 0$ and drops out of the habitat constraint. The parcels available for PES in period 2 come from two sources. Every parcel under easement is held in conservation for certain: C writes N^E easements and keeps all N^E of those parcels. Each of the remaining $N - N^E$ unprotected parcels survives only if its development offer falls short of working-land profit, which occurs with probability $1-\theta$.

[7] This is the natural multi-parcel analogue of the single-parcel benefit function and is what makes the allocation problem tractable; it specializes the per-parcel benefit to $B(x) = x$.

The available pool is therefore

$$P(N^E) = N^E + (1-\theta)(N - N^E) = (1-\theta)N + \theta\, N^E, \tag{12}$$

where N^E is the number of easements. The second expression collects terms in N^E and is the form we differentiate below: $dP/dN^E = \theta$, because the marginal easement adds a parcel to the pool only in the event—of probability θ—that the parcel would otherwise have developed. Under the uniform specialization, $P(N^E) = (\pi^L/\bar{D})N + [(\bar{D} - \pi^L)/\bar{D}]\, N^E$.

The conservationist's instruments are thus the number of easements N^E and the per-parcel PES level on each surviving parcel. We characterize the optimal PES allocation first, then reduce the whole problem to a single choice variable.

4.2 PES Spreading

Proposition 3. *Given P available parcels and a habitat target $\bar{B}$, total PES cost is minimized by setting the same input level $x_i = \bar{B}/P$ on every parcel.* (Proof: Appendix B.)

Because the parcels are identical and γ is strictly convex, equalizing the input across parcels is optimal: the first-order conditions for the cost-minimizing allocation require $\gamma'(x_i)$ to be equalized across parcels, which under identical γ forces x_i to be constant. Equivalently, by Jensen's inequality,

$$\gamma\left(\frac{1}{P}\sum_{i=1}^{P} x_i\right) \leq \frac{1}{P}\sum_{i=1}^{P} \gamma(x_i), \tag{13}$$

with equality only when x_i is constant, so spreading the target evenly minimizes total cost $P \cdot \gamma(\bar{B}/P)$. Due to the the convexity of γ, it is cheaper to ask many landowners to make small adjustments than to ask a few to make large ones. This result is central to landscape model because it collapses the multi-dimensional PES allocation to a function of a single quantity—the number of available parcels $P(N^E)$—and hence to a function of the single choice variable N^E.

4.3 The Cost-Minimization Problem

Using Proposition 3 to substitute $x_i = \bar{B}/P(N^E)$, total cost reduces to a function of N^E alone:

$$\min_{0 \leq N^E \leq N} \quad TC(N^E) = N^E \cdot C_E \;+\; \beta\, P(N^E) \cdot \gamma\left(\frac{\bar{B}}{P(N^E)}\right). \tag{14}$$

The first term is total easement spending; the second is total PES spending, discounted, with the per-parcel input set by spreading. This one-variable formulation replaces the case-by-case enumeration of Kuhn–Tucker conditions that a direct analysis on the multi-parcel problem would require.

First-order condition. Writing $x^* \equiv \bar{B}/P(N^E)$ for the common per-parcel input the target now requires across the available pool—in general different from the unconstrained single-parcel optimum x_2^P of Section 3—and differentiating, the marginal PES term simplifies, through the product rule and $dP/dN^E = \theta$, to $\theta\big[\gamma(x^*) - x^*\gamma'(x^*)\big]$ (the algebra is in Appendix B), so the optimum satisfies

$$C_E = \beta\theta\big[x^*\gamma'(x^*) - \gamma(x^*)\big] \equiv \beta\theta\, h(x^*), \tag{15}$$

where $h(x^*) \equiv x^*\gamma'(x^*) - \gamma(x^*)$ is the reduction in total PES spending from adding one parcel to the available pool, holding the habitat target fixed. Since $h(x^L) = 0$ and $h'(x^*) = x^*\gamma''(x^*) > 0$, the marginal value of an easement is positive whenever PES is active ($x^* > x^L$), so an interior optimum exists when development pressure is large enough for the left-hand side to be met.

A distribution-free optimum. Dividing (15) by θ and using $C_E = \theta\,\mathbb{E}[d - \pi^L \mid d > \pi^L]$ from Equation (5) yields

$$\mathbb{E}\big[d - \pi^L \mid d > \pi^L\big] = \beta\, h(x^*). \tag{16}$$

The development probability θ cancels entirely, and with it the shape of the offer distribution. The intuition is that θ scales both sides of the easement trade-off: an easement costs the option value of development, which is incurred only when development would have occurred (probability θ), and it saves PES costs only by preserving a parcel that would otherwise have developed (again probability θ). What remains is a balance between the conditional expected development surplus the landowner forgoes and the PES cost savings from one more parcel. Two landscapes with entirely different development-offer distributions but the same conditional mean surplus have the same optimal number of easements; the conservationist needs to know only that conditional mean, not the full distribution. Under the uniform specialization, Equation (16) becomes

$$\frac{\bar{D} - \pi^L}{2} = \beta\, h\!\left(\frac{\bar{B}}{P(N^E)}\right), \tag{17}$$

a single equation in N^E.

4.4 Easements Are Never Used Alone

The first-order condition (15) has an implication that separates the frictionless landscape from conservation practice, where pure easement programs are common.

Proposition 4. *In the model without transaction costs, it is never optimal to purchase easements without also conducting PES. At any optimum with $N^E > 0$, the conservationist sets $x_i > x^L$ on the available parcels.* (Proof: Appendix B.)

Suppose to the contrary that C bought easements and set $x_i = x^L$, meeting the target through baseline land use alone. The marginal benefit of an easement is the PES cost it saves—the right-hand side of Equation (15), $\beta\theta\, h(x^*)$—and at $x^* = x^L$ this is exactly zero, because $\gamma(x^L) = 0$ and $\gamma'(x^L) = 0$ by Lemma 1. But the marginal cost of that easement is $C_E > 0$ whenever development is possible. An easement that merely pushes the landscape to $x_i = x^L$ therefore costs a positive amount and saves nothing, so C always prefers a marginal amount of PES—which is free at x^L—to the last such easement. Pure easement programs are thus a puzzle for the frictionless model; Section 5 shows that per-contract monitoring costs resolve it.

4.5 Comparative Statics

The landscape model's central result is its comparative static on development pressure, which overturns the single-parcel result. At an interior optimum, define $\Phi(N^E; \bar{D}, \bar{B}, N) \equiv (\bar{D} - \pi^L)/2 - \beta\, h(\bar{B}/P(N^E))$, the difference between the two sides of (17). Since $h' > 0$ and P is increasing in N^E, we have $\partial\Phi/\partial N^E > 0$, so by the implicit function theorem the sign of $dN^E/d\alpha$ matches the sign of $-\partial\Phi/\partial\alpha$ for any parameter α.

Development pressure ($\bar{D}$): ambiguous. Higher development pressure has two opposing effects on the optimal number of easements (Appendix B derives the expression):

$$\frac{\partial\Phi}{\partial\bar{D}} = \frac{1}{2} - \beta\, x^*\gamma''(x^*) \cdot \frac{\bar{B}\pi^L(N - N^E)}{\bar{D}^2 P^2}. \tag{18}$$

The first term is a *cost* channel: a higher $\bar{D}$ raises the conditional development surplus $(\bar{D} - \pi^L)/2$ and makes easements costlier. The second is a *spreading* channel: a higher $\bar{D}$ raises the development probability θ, shrinks the pool of surviving parcels P, drives up per-parcel PES costs through convexity, and so makes each easement more valuable. The net effect depends on which channel dominates. When development pressure is mild, the spreading channel dominates and easements rise with pressure; when development pressure is severe, the cost channel dominates and easements fall. The optimal number of easements can therefore be non-monotonic in development pressure, rising and then falling—an inverted-U. Endogenizing development thus overturns the single-parcel intuition, where a higher $\bar{D}$ unambiguously discouraged easements because the single parcel offered no spreading margin to offset the higher cost.

Habitat target ($\bar{B}$): positive; total parcels (N): negative. A more ambitious target raises x^* and hence the marginal value of preserving parcels, so $dN^E/d\bar{B} > 0$. A larger landscape provides more parcels across which to spread PES even without easements, lowering the marginal value of each easement, so $dN^E/dN < 0$.

Working-land profit (π^L): ambiguous. Higher working-land profit lowers both the development probability and the easement cost—which makes easements cheaper and tends to raise N^E—but it also raises the number of parcels surviving without protection, which lowers the marginal value of easements and tends to reduce N^E. As with development pressure, the net effect depends on the balance of the two channels. Table 1 collects the comparative statics.

Table 1. Comparative Statics for the Optimal Number of Easements (Uniform Case)

Parameter	$dN^E/d(\cdot)$	Mechanism
$\bar{D}$ (development pressure)	ambiguous	Raises both the cost and the value of easements (possible inverted-U)
$\bar{B}$ (habitat target)	$+$	A harder target requires more parcels for PES spreading
N (total parcels)	$-$	More parcels to spread PES across, even without easements
π^L (working-land profit)	ambiguous	Lowers θ and C_E, but also raises the surviving pool P

The ambiguity of development pressure is the model's central comparative static. In a formulation with a free probability θ and a free development value, the two would deliver opposing but separable results—a higher value discourages easements, a higher probability encourages them. Tying both to a single primitive $\bar{D}$ shows that they operate at once: development pressure raises the cost of easements and the value of the parcels they preserve in the same breath, and which force wins is an empirical feature of the landscape rather than a foregone conclusion.

4.6 Three Regimes

The comparative statics describe how an interior optimum moves; we now characterize when the optimum is interior at all. The solution to (14) takes one of three forms, depending on how the target compares to what baseline land use and PES can deliver. We state the boundaries under the uniform specialization, where $1 - \theta = \pi^L/\bar{D}$; Appendix B derives each one.

1. **Do nothing** ($N^E = 0$, no PES). Baseline land use on surviving parcels already meets the target: $(\pi^L/\bar{D})\, N\, x^L \geq \bar{B}$, i.e., $\bar{D} \leq \pi^L N x^L/\bar{B}$.

2. **PES only** ($N^E = 0$, PES on surviving parcels). The target is not met by baseline land use, but the marginal benefit of the first easement falls short of its cost:

$$\frac{\bar{D} - \pi^L}{2} > \beta\, h\left(\frac{\bar{B}\bar{D}}{\pi^L N}\right). \tag{19}$$

3. **PES plus easements** (interior $0 < N^E < N$, or the corner $N^E = N$). The first easement's marginal benefit exceeds its cost, and N^E is set by the first-order condition (17), reaching the corner $N^E = N$ when $(\bar{D} - \pi^L)/2 \leq \beta\, h(\bar{B}/N)$.

Figure 2 maps these regimes in the space of development pressure $\bar{D}$ and the habitat target $\bar{B}$. The "do nothing" boundary is the hyperbola $\bar{B} = (\pi^L/\bar{D})N x^L$, above which C must act; the "PES only" region sits above it at moderate development pressure; and easements enter only once development pressure is high enough that preserving parcels for spreading is worth the option value forgone—the spreading channel of Equation (18) at work.

Figure 2. Optimal Conservation-Instrument Regimes in $(\bar{D}, \bar{B})$ Space

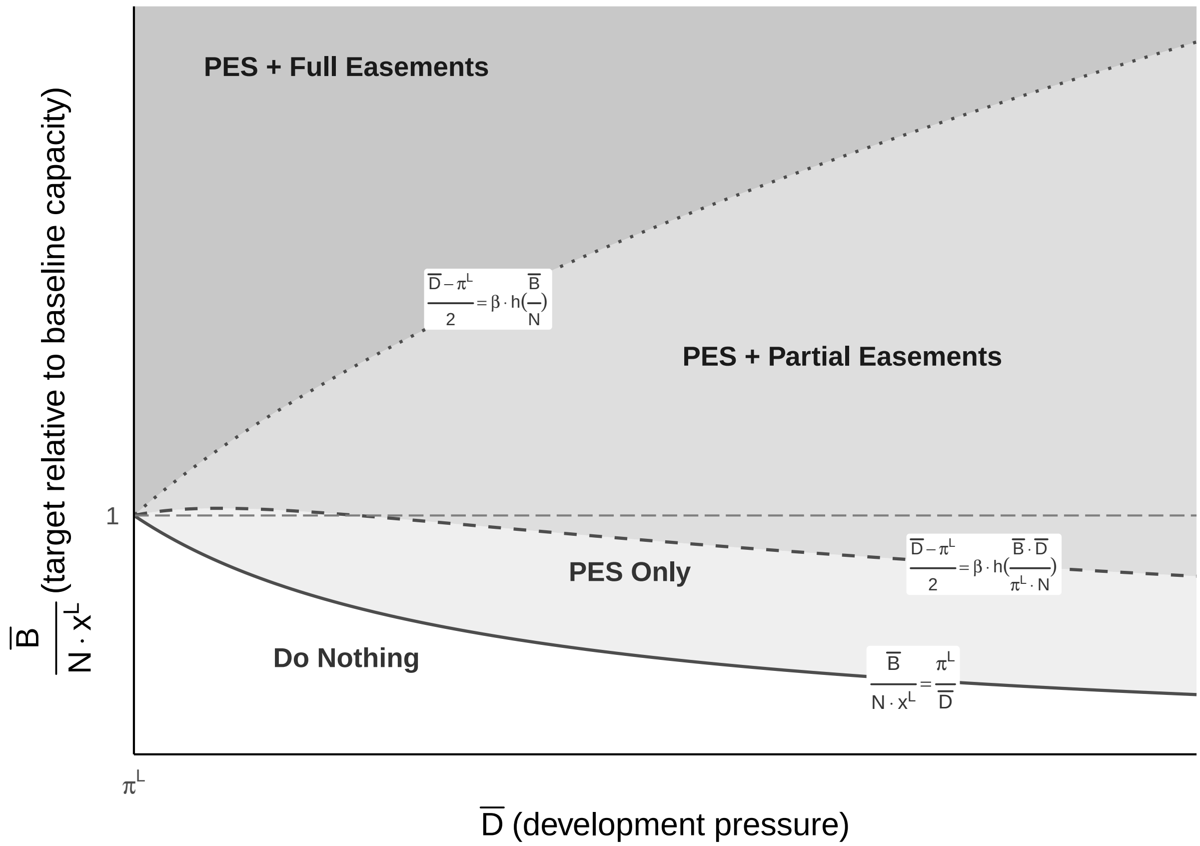

Notes: The solid curve is the "do nothing" boundary $\bar{B} = (\pi^L/\bar{D})Nx^L$; the dashed curve is the first-order condition evaluated at $N^E = 0$, separating "PES only" from "PES plus easements"; the dotted curve is the first-order condition at $N^E = N$, above which all parcels receive easements. Plotted under the uniform specialization with a quadratic PES cost function.

5 Extensions

The landscape model assumes identical parcels and frictionless contracting, and we relax each assumption in this section. Heterogeneity in landowners delivers a targeting logic that maps instruments to parcel types and rationalizes the observed coexistence of PES and easements across a landscape. Per-contract monitoring costs rehabilitate the two instruments the frictionless model rules out—fee-simple acquisition and pure easements—and tie their resurgence to the cost of administering incentive contracts.

5.1 Heterogeneous Landowners

Suppose the landscape contains two types of parcels, A and B, with n_A and n_B parcels respectively and $n_A + n_B = N$. Types may differ in their baseline land use x_j^L, their development pressure $\bar{D}_j$ (and hence development probability θ_j and easement cost $C_{E,j}$), and the curvature of their

PES cost function γ_j. The parameterization through a single primitive $\bar{D}_j$ continues to generate both the probability and the cost of development for each type, so a type facing more intense development pressure has both a higher development probability and a higher easement cost.

Proposition 3 applies within each type, so C spreads PES evenly across all available parcels of a given type, $x_j = \bar{B}_j/P_j$, where $\bar{B}_j$ is the share of the target allocated to type j and the available pool is, as in Equation (12), the eased parcels plus the unprotected survivors: $P_j(N_j^E) = N_j^E + (1-\theta_j)(n_j - N_j^E)$. The new margin is the allocation of the target across types, alongside the type-specific easement choices. The conservationist solves

$$\min_{N_A^E,\, N_B^E,\, \bar{B}_A} \quad \sum_{j\in\{A,B\}} \left[N_j^E \cdot C_{E,j} + \beta\, P_j \cdot \gamma_j\left(\frac{\bar{B}_j}{P_j}\right)\right] \tag{20}$$

subject to $\bar{B}_A + \bar{B}_B = \bar{B}$, the easement bounds $0 \leq N_j^E \leq n_j$, and $\bar{B}_j/P_j \geq x_j^L$. Two first-order conditions characterize the solution. Allocating the target across types equalizes the marginal PES cost,

$$\gamma_A'(x_A^*) = \gamma_B'(x_B^*), \tag{21}$$

a standard equimarginal condition, while easements on each type satisfy the same balance as in the homogeneous model,

$$C_{E,j} = \beta\theta_j\, h_j(x_j^*), \qquad x_j^* = \bar{B}_j/P_j(N_j^E). \tag{22}$$

These conditions yield a targeting logic for each instrument. The equimarginal condition (21) directs more PES toward the type with the lower marginal PES cost—the type whose baseline land use is closer to the conservation goal, or whose γ_j is flatter.

Proposition 5. *C directs more PES toward parcel types whose marginal cost of conservation is lower, that is, types whose baseline land use is closer to the conservation objective.* (Proof: Appendix B.)

This matches the BirdReturns intuition: flood-irrigated rice fields require only a small input adjustment to provide migratory habitat, so they absorb a large share of the conservation effort at low cost. Easement placement follows from comparing (22) across types.

Proposition 6. *C places easements preferentially on parcel types where the easement is cheap (low development pressure $\bar{D}_j$) or where the PES cost saved by preserving a parcel, $\beta\theta_j\, h_j(x_j^*)$, is large (steep γ_j or intensive per-parcel PES).* (Proof: Appendix B.)

The two propositions together generate a sorting of instruments across the landscape. The relevant case for intensive agriculture is one in which development pressure and PES cost are positively correlated, so that the parcels most exposed to development are also the most expensive to reach through incentive payments. There the conservationist faces a tension—such parcels are costly to protect by easement and costly to adjust by PES—and the optimal mix depends on the strength of the correlation. Figure 3 maps the assignment of instruments over the two dimensions of heterogeneity: parcel types at low development pressure receive PES alone; as development pressure or the steepness of the PES cost function rises, easements enter the mix, and the share of parcels eased grows with both.

Figure 3. Instrument Assignment Across Parcel Types

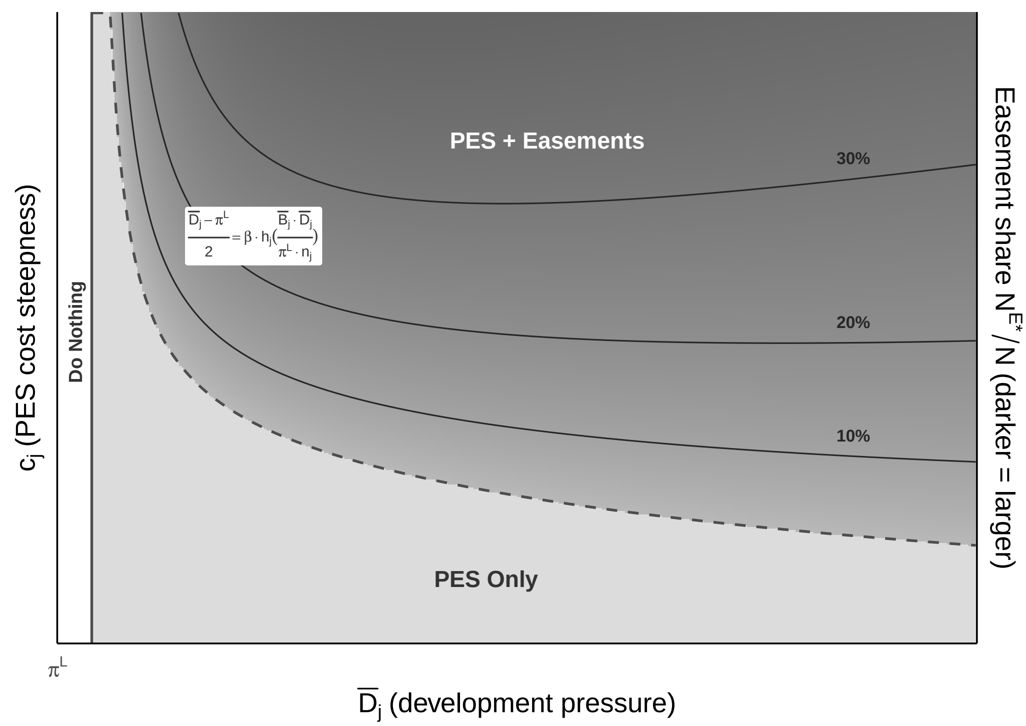

Notes: Assignment is shown in the space of development pressure $\bar{D}_j$ and PES cost steepness c_j (with $\gamma_j(x) = c_j(x - x^L)^2$). The dashed curve is the type-level easement condition—the first easement's marginal benefit equals its cost at $N_j^E = 0$—above which easements enter the optimal mix; shading within that region gives the optimal share of parcels eased. Plotted under the uniform specialization with the target fixed at $\bar{B}_j = 0.9\, n_j x^L$.

5.2 Monitoring and Transaction Costs

The frictionless model rules out two instruments that are common in practice—fee-simple acquisition (Proposition 1) and pure easement programs (Proposition 4). Both results turn on PES being costless to administer. We now show that a per-contract monitoring cost restores both instruments, and we characterize the threshold at which it does so.

Every PES contract requires C to verify the landowner's compliance, which we model as a fixed cost $M > 0$ incurred on each parcel under PES, regardless of the input level. Monitoring costs include search, contracting, verification, and enforcement. On a single parcel, the PES cost of achieving x_2 becomes $\gamma(x_2) + M$ while acquisition remains $\pi^L + w_2 x_2$, so Proposition 1 can fail for large M. At input levels far from x^L—where the conservationist wishes to suppress the conservation-relevant input substantially—the fixed monitoring cost can push the PES schedule above the acquisition line, and acquisition becomes the low-cost way to secure that input. This accounts for why ownership persists: it is most attractive precisely where C wishes to depart most from the landowner's optimum, so that the avoided monitoring cost outweighs the purchase price.

At the landscape scale, monitoring costs make the total cost function piecewise. Define the *easement-only threshold* N_0^E as the number of easements at which baseline land use on surviving parcels exactly meets the target, $P(N_0^E) \cdot x^L = \bar{B}$; this exists in $(0, N]$ when $(1-\theta)Nx^L < \bar{B} \leq Nx^L$. Total cost is then

$$TC(N^E) = \begin{cases} N^E \cdot C_E + \beta\, P(N^E)\big[\gamma(\bar{B}/P(N^E)) + M\big] & N^E < N_0^E \quad \text{(PES active)}, \\ N^E \cdot C_E & N^E \geq N_0^E \quad \text{(easement only)}. \end{cases} \tag{23}$$

In the PES-active region, every contract bears the fixed cost M; in the easement-only region, no PES is written and no monitoring cost is incurred. The switch between the two creates a discontinuity.

Proposition 7. *For $M > 0$, total cost jumps down at the easement-only threshold:*

$$\lim_{N^E \to N_0^{E-}} TC(N^E) - TC(N_0^E) = \beta\, \frac{\bar{B}}{x^L}\, M > 0. \tag{24}$$

(Proof: Appendix B.)

As N^E approaches N_0^E from below, the per-parcel input approaches x^L and the PES payment $\gamma(x^*)$ vanishes by Lemma 1—but the monitoring cost M does not, so the conservationist is paying the full monitoring bill $\beta(\bar{B}/x^L)M$ to administer contracts that ask landowners for a vanishingly small adjustment. The last easement, which pushes the landscape to baseline land use and eliminates the need for any PES, saves that entire bill at a stroke. The frictionless model omits this margin. Without monitoring costs, the zero marginal cost of PES at x^L (Proposition 4) always makes a little PES preferable to the last easement; with them, the fixed cost of administering even trivial PES tips the balance.

In the PES-active region the first-order condition becomes

$$C_E = \beta\theta\big[h(x^*) - M\big] \tag{25}$$

(the derivation parallels that of Equation (15) in Appendix B, with per-parcel cost $\gamma(x^*) + M$ in place of $\gamma(x^*)$), so monitoring lowers the marginal value of each easement by $\beta\theta M$: an easement preserves a parcel that must then be monitored. An interior solution with both instruments requires the per-parcel spreading saving to exceed the monitoring cost, $h(x^*) > M$. Within this region, more costly monitoring leads to *fewer* easements, as C economizes on monitoring by allowing more parcels to develop and conducting more intensive PES on fewer survivors. But this interior logic is overturned once monitoring is costly enough to trigger the easement-only corner.

Proposition 8. *For any configuration with $0 < N_0^E \leq N$, there is a threshold $M^* > 0$ such that the globally optimal strategy involves PES for $M < M^*$ and is easement-only—$N^E = N_0^E$ with $x_i = x^L$ on every parcel—for $M \geq M^*$.* (Proof: Appendix B.)

The easement-only cost $N_0^E \cdot C_E$ does not depend on M, while the optimized PES-inclusive cost rises in M at rate βP^* by the envelope theorem. At $M = 0$ the PES-inclusive strategy is cheaper (Proposition 4); as M grows the two cross at a unique M^*, above which paying for enough easements to meet the target through baseline land use alone is cheaper than paying to monitor any PES. The optimal number of easements is therefore non-monotonic in M: it falls as monitoring costs rise within the interior region, then jumps up to N_0^E at M^*. Figure 4 shows the easement-only region that monitoring opens up, at moderate targets and high development pressure.

The easement-only result gives a formal rationale for the pure easement programs that the frictionless model cannot explain—the USDA Agricultural Conservation Easement Program and state farmland-preservation programs that protect land without writing incentive contracts. It also reframes the historical trajectory of conservation. The long shift from fee-simple acquisition toward easements and then toward PES coincides with a steady decline in monitoring costs, as crowd-sourced and satellite data have made compliance cheap to verify: BirdReturns relies on smartphone-based bird counts and remote sensing to know when and where to pay for flooding and to confirm it occurs. In the model, a falling M collapses the easement-only region back toward

Figure 4. Optimal Instrument Regimes with Monitoring Costs

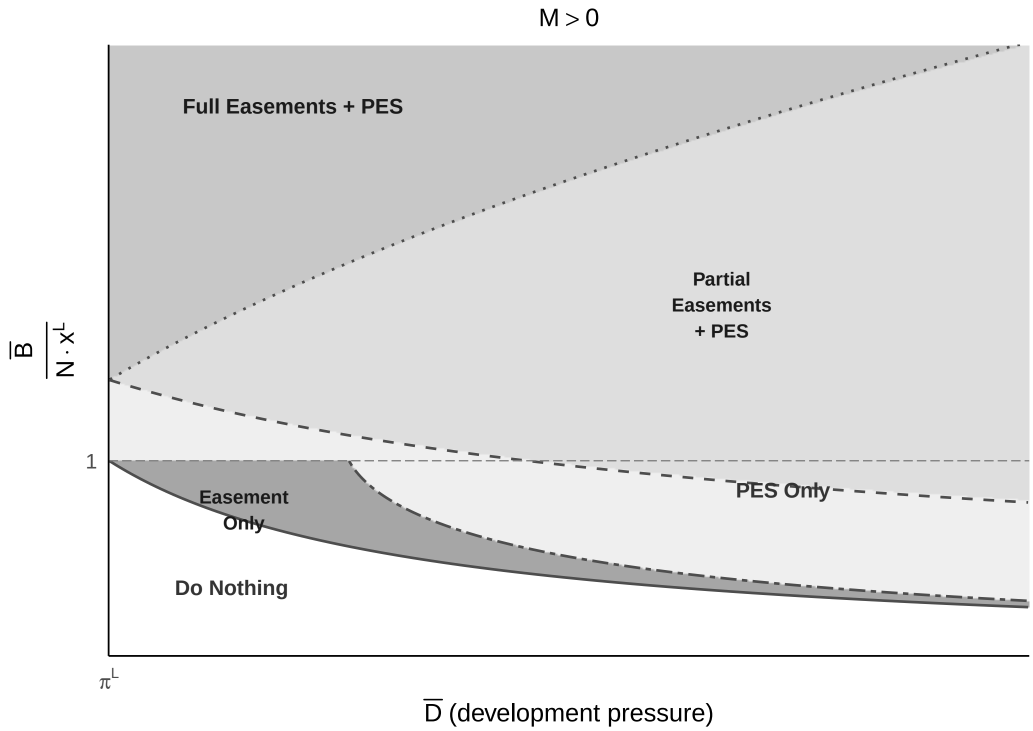

Notes: Regimes are shown in $(\bar{D},\ \bar{B}/(Nx^L))$ space with monitoring costs ($M > 0$). Relative to the frictionless baseline of Figure 2, monitoring opens an *easement-only* region (at moderate targets and high development pressure) where C meets the target through easements alone to avoid per-contract monitoring, and the PES-inclusive boundaries shift upward as monitoring raises the cost of PES.

the frictionless baseline and expands the region where PES is part of the optimal mix. Because the regime boundaries shift discontinuously with M, even modest improvements in monitoring technology can trigger discrete changes in the optimal instrument—an account of why PES has proliferated as monitoring has become cheap.

Heterogeneous monitoring costs would add a further layer of targeting: PES on parcels whose compliance is cheap to verify, such as row crops with satellite-observable practices, and easements or acquisition on parcels whose management is hard to monitor. This sharpens the assignment of Figure 3 along a third dimension, which we leave to future work.

6 Discussion

The model abstracts from several features of real conservation landscapes. We discuss four that shape the instrument choice—overlapping policies, spatial connectivity, the flexibility of temporary contracts, and an empirical mapping of the model's regimes—and indicate how each would extend the analysis.

6.1 Overlapping Policies

Conservation contracts are written against a backdrop of other policies that shape the model's primitives. Zoning that limits development truncates the upper tail of the development-offer distribution, lowering $\bar{D}$ and so reducing both the probability of development θ and the easement cost C_E; where zoning already constrains conversion, easements are cheaper but also less necessary. Tax treatment matters as well. Many counties tax working land at reduced rates, which lowers the return to development and makes easements cheaper, and conservation easements often carry their own tax benefits that reduce their effective cost—Parker and Thurman (2018) find that more generous incentives induce additional easements. Agricultural subsidies raise working-land profit π^L and therefore move both the easement cost and the PES cost. These policies do not change the structure of the instrument choice, but they shift the regime boundaries of Figures 2 and 4 and should be accounted for in any application.

6.2 Connectivity and Spatial Structure

The landscape model treats parcels as interchangeable, which is the assumption that delivers the spreading result but also sets aside the spatial configuration that matters for many species. Where habitat value depends on connectivity—corridors, stepping stones, contiguous core areas—the habitat function is no longer additive across parcels, and losing a single parcel in a corridor can cause a discontinuous drop in value. Connectivity makes easements relatively more valuable, because PES leaves each parcel exposed to development in the next period and a single conversion can break a corridor, whereas an easement guarantees the parcel's place in the landscape. This is also the one setting in which a pure easement strategy can be optimal absent monitoring costs: for a critical corridor parcel, the non-convexity in habitat value can outweigh the spreading logic of Proposition 4. The pattern is consistent with the concentration of easements along migration corridors such as the Path of the Pronghorn in the GYE, alongside the use of PES on the more interchangeable stopover habitat of the Central Valley. A tractable spatial model—parcels arrayed on a line with a corridor requiring a run of consecutive protected parcels—would turn the problem from how many parcels to protect into which ones, and we leave it to future work.

6.3 The Flexibility of Temporary Contracts

The model gives PES a cost advantage but not the flexibility advantage emphasized elsewhere (Reynolds et al., 2017). When the conservation value of particular parcels changes over time, temporary contracts let the conservationist re-allocate effort as conditions change, while easements bind permanently. This flexibility favors PES where habitat needs are shifting or uncertain, and it can also make acquisition relatively more attractive than easements, since land can be sold if its conservation value falls whereas an easement cannot be undone. Incorporating stochastic, time-varying benefits would add this margin to the instrument comparison.

7 Conclusion

Conservation organizations and government agencies have come to rely on a widening menu of contractual arrangements for protecting private working land, yet economics has offered little guidance on how to choose among them. We place fee-simple acquisition, conservation easements, and payments for ecosystem services in a single cost-minimization problem and characterize when each belongs in the optimal mix.

The baseline results can be stated simply. On a single parcel, it is never cheaper to buy land outright than to pay the landowner to provide conservation, because the landowner continues to earn profit from the parcel; and whether to add an easement depends on the conservation value of preventing development relative to the landowner's gain from developing, not on how likely development is. At the landscape scale, the convexity of incentive costs leads the conservationist to spread payments thinly across many parcels, and easements acquire a second rationale beyond preventing development—preserving parcels over which to spread that effort. The result is that development pressure cuts both ways, raising the cost and the value of easements together, so its net effect on the optimal number of easements is ambiguous and may be non-monotonic.

These results carry a clear implication for the emphasis on public acquisition in the 30 by 30 agenda. In the frictionless model, acquisition is dominated and pure easement programs are never optimal, so a conservation strategy built on buying land is difficult to justify on cost grounds alone. That conclusion is not the end of the argument, however, because the same model shows what would rehabilitate ownership and easements: the cost of administering incentive contracts. When monitoring is expensive—particularly where the conservationist wishes to suppress inputs substantially, or where compliance is hard to verify—acquisition and pure easements re-emerge as the least-cost options. The historical trajectory from acquisition toward easements and then toward leasing tracks the falling cost of monitoring, as remote sensing and crowd-sourced data have made compliance cheap to observe, and the model predicts that continued declines will further favor PES. Advocates for a particular instrument, then, should be explicit about the monitoring environment they assume: the case for acquisition rests on frictions that are themselves receding.

To our knowledge this is the first analysis to treat acquisition, easements, and PES as alternative solutions to a common conservation problem rather than in isolation. In doing so it begins to supply for conservation what the theory of the firm supplies for production—an account of why an organization chooses ownership in some circumstances, partial property rights in others, and short-term contracts in still others. Much remains to be done. A spatial model would bring connectivity and corridors into the instrument choice; time-varying benefits would add the flexibility advantage of temporary contracts; and an empirical mapping of real landscapes into the model's regimes would test its predictions against observed conservation practice. We hope the framework developed here provides a useful starting point for that work.

A Notation Reference

Symbol	Definition
x_1, x_2	Landowner inputs (1 = conservation-neutral, 2 = conservation-relevant)
$x^L \equiv x_2^L$	Landowner's profit-maximizing choice of x_2
π^L	Landowner's maximum profit, $\pi(x_1^L, x_2^L)$
$B(x_2)$	Conservation benefit function (concave, $B(0) = 0$)
$\gamma(\bar{x}_2)$	PES cost function: L's willingness to accept $x_2 = \bar{x}_2$ (convex, min 0 at x^L)
x_2^P	Optimal PES input level on a single parcel
λ	Marginal value of conservation benefit (multiplier on the target)
d	Development offer received by a parcel, $d \sim F(d)$ on $[0, \bar{D}]$
$F(d)$	CDF of the development-offer distribution
$\bar{D}$	Upper bound of the offer distribution (development pressure)
θ	Development probability, $1 - F(\pi^L)$; uniform: $(\bar{D} - \pi^L)/\bar{D}$
C_E	Easement cost per parcel, $\mathbb{E}[\max(d - \pi^L, 0)]$; uniform: $(\bar{D} - \pi^L)^2/(2\bar{D})$
β	Discount factor
N	Total number of parcels in the landscape
N^E	Number of easements purchased
$P(N^E)$	Parcels available in period 2, $(1 - \theta)N + \theta N^E$
$\bar{B}$	Landscape habitat target
x^*	Per-parcel PES level, $\bar{B}/P(N^E)$
$h(x^*)$	PES cost savings from one more parcel, $x^*\gamma'(x^*) - \gamma(x^*)$
n_j	Number of type-j parcels (heterogeneous model)
$\bar{D}_j$	Development pressure for type j; yields θ_j and $C_{E,j}$
P_j	Available type-j parcels, $(1 - \theta_j)n_j + \theta_j N_j^E$
$\bar{B}_j$	Habitat target allocated to type j
M	Per-contract monitoring/transaction cost (Section 5.2)
N_0^E	Easement-only threshold: N^E such that $P(N_0^E)x^L = \bar{B}$
M^*	Monitoring-cost threshold above which easement-only is optimal

B Proofs

Proof of Lemma 1

Write the PES cost function explicitly as

$$\gamma(\bar{x}_2) = \big[pf(x_1^L, x_2^L) - w_1 x_1^L - w_2 x_2^L\big] - \big[pf(\tilde{x}_1, \bar{x}_2) - w_1\tilde{x}_1 - w_2\bar{x}_2\big], \tag{26}$$

where $\tilde{x}_1 = \tilde{x}_1(\bar{x}_2)$ is L's profit-maximizing choice of x_1 given $x_2 = \bar{x}_2$.

Minimum of zero at x^L. Since $\tilde{x}_1(x_2^L) = x_1^L$, the two bracketed terms coincide at $\bar{x}_2 = x_2^L$, so $\gamma(x_2^L) = 0$. Because profit at any $(x_1, x_2) \neq (x_1^L, x_2^L)$ is strictly below π^L, we have $\gamma(\bar{x}_2) \geq 0$ with equality only at x_2^L.

First derivative. Differentiating and applying the envelope theorem to L's choice of $\tilde{x}_1$,

$$\frac{d\gamma}{d\bar{x}_2} = -\left[\left(p\frac{\partial f(\tilde{x}_1,\bar{x}_2)}{\partial \bar{x}_2} - w_2\right) + \frac{d\tilde{x}_1}{d\bar{x}_2}\underbrace{\left(p\frac{\partial f(\tilde{x}_1,\bar{x}_2)}{\partial \tilde{x}_1} - w_1\right)}_{=0 \text{ by } L\text{'s FOC}}\right] = w_2 - p\frac{\partial f(\tilde{x}_1,\bar{x}_2)}{\partial \bar{x}_2}. \tag{27}$$

The underbraced term is L's first-order condition for x_1 under the constraint $x_2 = \bar{x}_2$ and so equals zero. At $\bar{x}_2 = x_2^L$, Equation (27) is zero by L's unconstrained first-order condition for x_2. For $\bar{x}_2 < x_2^L$, concavity of f gives $p\,\partial f/\partial \bar{x}_2 > w_2$, so $\gamma' < 0$; for $\bar{x}_2 > x_2^L$, $\gamma' > 0$. Thus γ decreases to x_2^L and increases thereafter.

Convexity. Differentiating the simplified first derivative $\gamma' = w_2 - p\,f_2(\tilde{x}_1(\bar{x}_2), \bar{x}_2)$ from Equation (27),

$$\frac{d^2\gamma}{d\bar{x}_2^2} = -p\left[f_{22} + f_{21}\,\frac{d\tilde{x}_1}{d\bar{x}_2}\right]. \tag{28}$$

The constrained first-order condition $p\,f_1(\tilde{x}_1, \bar{x}_2) = w_1$ implies $\frac{d\tilde{x}_1}{d\bar{x}_2} = -f_{12}/f_{11}$, so

$$\frac{d^2\gamma}{d\bar{x}_2^2} = -p\,\frac{f_{11}f_{22} - f_{12}^2}{f_{11}} > 0, \tag{29}$$

which is positive because $f_{11} < 0$ and $f_{11}f_{22} - f_{12}^2 > 0$ by concavity of f.

Slope bound. From Equation (27), $\gamma' = w_2 - p\,\partial f/\partial \bar{x}_2$; since $\partial f/\partial \bar{x}_2 \geq 0$, we have $\gamma' \leq w_2$ for all $\bar{x}_2$. This upper bound is what the dominance result in Proposition 1 requires. □

Proof of Proposition 1

Acquisition costs $\pi^L + w_2 x_2$. Under PES, L always retains the option of shutting down conventional production and supplying x_2 alone, which costs C at most $\pi^L + w_2 x_2$; hence $\gamma(0) \leq \pi^L$. Together with the slope bound $\gamma' \leq w_2$ from Lemma 1, integrating gives $\gamma(x_2) \leq \pi^L + w_2 x_2$ for all $x_2 \geq 0$. PES therefore secures any input level at weakly lower cost; allowing C to choose a different input level under PES makes the advantage generically strict. In the two-period model, the expected cost of the development option, βC_E, enters acquisition and easement-plus-PES identically and so cancels from the comparison. □

Proof of Proposition 2

Value conservation on the parcel at the optimal PES contract, which yields discounted net benefit $\beta[B(x_2^P) - \gamma(x_2^P)]$ in any period-2 state in which the parcel is conserved. Without an easement, the parcel is conserved only if it does not develop, which occurs with probability $1-\theta$, so the expected payoff is $\beta(1-\theta)[B(x_2^P) - \gamma(x_2^P)]$. With an easement, C pays C_E in period 1 and conserves for certain, for a payoff of $\beta[B(x_2^P) - \gamma(x_2^P)] - C_E$. The easement is preferred when

$$\beta\theta\big[B(x_2^P) - \gamma(x_2^P)\big] > C_E = \theta\,\mathbb{E}[d - \pi^L \mid d > \pi^L]. \tag{30}$$

Dividing by $\theta > 0$ and substituting $\gamma(x_2^P) = \pi^L - \pi(\tilde{x}_1, x_2^P)$ gives condition (11), in which θ does not appear. □

Proof of Proposition 3

C minimizes $\sum_{i=1}^P \gamma(x_i)$ subject to $\sum_i x_i \geq \bar{B}$ and $x_i \geq x^L$. When $\bar{B}/P > x^L$ the lower bounds slacken and each interior first-order condition reads $\gamma'(x_i) = \lambda$; since γ is strictly convex, γ' is strictly increasing and invertible, so $x_i = (\gamma')^{-1}(\lambda)$ is the same for every i, giving $x_i = \bar{B}/P$. Equivalently, strict convexity and Jensen's inequality give $\gamma(\frac{1}{P}\sum_i x_i) \leq \frac{1}{P}\sum_i \gamma(x_i)$ with equality only at constant x_i, so equal spreading minimizes $P\,\gamma(\bar{B}/P)$. □

Derivations for Section 4

The first-order condition (15). Differentiating the total cost function (14) with respect to N^E, using $dP/dN^E = \theta$ from Equation (12) and writing $x^* = \bar{B}/P(N^E)$,

$$\frac{d}{dN^E}\left[\beta\, P \cdot \gamma\!\left(\frac{\bar{B}}{P}\right)\right] = \beta\left[\theta\,\gamma(x^*) + P\,\gamma'(x^*)\cdot\left(-\frac{\bar{B}}{P^2}\right)\theta\right] = \beta\theta\big[\gamma(x^*) - x^*\gamma'(x^*)\big], \tag{31}$$

where the first term in the bracket is the cost of writing PES on the marginal preserved parcel and the second is the saving from lowering the input on every inframarginal parcel. Hence $dTC/dN^E = C_E - \beta\theta\big[x^*\gamma'(x^*) - \gamma(x^*)\big]$, and setting this to zero gives Equation (15).

The comparative-static expression (18). Under the uniform specialization, $\theta = 1 - \pi^L/\bar{D}$, so $\partial\theta/\partial\bar{D} = \pi^L/\bar{D}^2$. Writing $P = N - \theta(N - N^E)$ from Equation (12),

$$\frac{\partial P}{\partial \bar{D}} = -(N - N^E)\,\frac{\pi^L}{\bar{D}^2}, \qquad \frac{\partial x^*}{\partial \bar{D}} = -\frac{\bar{B}}{P^2}\,\frac{\partial P}{\partial \bar{D}} = \frac{\bar{B}\,\pi^L\,(N - N^E)}{\bar{D}^2 P^2}. \tag{32}$$

Since $h'(x^*) = \gamma'(x^*) + x^*\gamma''(x^*) - \gamma'(x^*) = x^*\gamma''(x^*)$, differentiating $\Phi = (\bar{D} - \pi^L)/2 - \beta\, h(x^*)$ gives

$$\frac{\partial \Phi}{\partial \bar{D}} = \frac{1}{2} - \beta\, x^*\gamma''(x^*)\cdot\frac{\bar{B}\,\pi^L\,(N - N^E)}{\bar{D}^2 P^2}, \tag{33}$$

which is Equation (18).

The regime boundaries. With $N^E = 0$, the surviving pool is $P(0) = (1-\theta)N = (\pi^L/\bar{D})N$ parcels, each supplying x^L at baseline, so doing nothing meets the target if and only if $(\pi^L/\bar{D})Nx^L \geq \bar{B}$, i.e., $\bar{D} \leq \pi^L N x^L/\bar{B}$. When that fails, $N^E = 0$ remains optimal if and only if the first easement's marginal benefit falls short of its cost, $dTC/dN^E \geq 0$ at $N^E = 0$; dividing by θ as in Equation (16) and substituting $x^* = \bar{B}/P(0) = \bar{B}\bar{D}/(\pi^L N)$ gives the PES-only condition $(\bar{D} - \pi^L)/2 > \beta\, h\big(\bar{B}\bar{D}/(\pi^L N)\big)$. Otherwise easements are used, with N^E set by Equation (17); the corner $N^E = N$ obtains when the marginal benefit still exceeds the cost at $N^E = N$, where $P(N) = N$ and $x^* = \bar{B}/N$, i.e., when $(\bar{D} - \pi^L)/2 \leq \beta\, h(\bar{B}/N)$. □

Proof of Proposition 4

Suppose C writes $N^E > 0$ easements and sets $x_i = x^L$. The marginal condition for N^E is $C_E = \beta\theta[x^L\gamma'(x^L) - \gamma(x^L)]$. By Lemma 1, $\gamma(x^L) = 0$ and $\gamma'(x^L) = 0$, so the right-hand side is zero while the left-hand side is $C_E > 0$ whenever $\bar{D} > \pi^L$. The marginal cost of an easement is thus strictly positive while its marginal benefit at $x_i = x^L$ is zero, so C strictly prefers a marginal amount of PES—costless at x^L—to the last easement. Hence any optimum with $N^E > 0$ has $x_i > x^L$. □

Proofs of Propositions 5 and 6

Proposition 5. The allocation of the target across types satisfies the equimarginal condition (21), $\gamma_A'(x_A^*) = \gamma_B'(x_B^*) = \lambda$. Each γ_j is strictly convex by Lemma 1, so γ_j' is strictly increasing and invertible and $x_j^* = (\gamma_j')^{-1}(\lambda)$. If type A has the lower marginal cost of conservation—$\gamma_A'(x) \leq \gamma_B'(x)$ for all x in the relevant range, with strict inequality at the optimum, as when type A's baseline use lies closer to the conservation objective or its cost function is flatter—then $(\gamma_A')^{-1}(\lambda) > (\gamma_B')^{-1}(\lambda)$, so $x_A^* > x_B^*$: the lower-marginal-cost type receives the larger per-parcel adjustment, and hence the larger share of the target per available parcel. □

Proposition 6. Easements on type j are governed by the type-j analogue of Equation (15): at an interior optimum $C_{E,j} = \beta\theta_j h_j(x_j^*)$, and type j receives easements at all ($N_j^E > 0$) only if the first easement's marginal benefit exceeds its cost, $\beta\theta_j h_j(x_j^*) \geq C_{E,j}$ at $N_j^E = 0$. Holding the PES side fixed, a lower $\bar{D}_j$ lowers the conditional development surplus and hence $C_{E,j}$, relaxing the condition; holding the easement cost fixed, a steeper γ_j or a larger per-parcel input x_j^* raises $h_j(x_j^*)$ (since $h_j' = x\gamma_j'' > 0$), also relaxing it. Easements therefore flow first to types where they are cheap or where the PES cost saved by preserving a parcel is large. □

Proof of Proposition 7

From Equation (23), as $N^E \to N_0^{E-}$ the per-parcel input $x^* \to x^{L+}$, so $\gamma(x^*) \to 0$ by Lemma 1. The per-parcel PES cost approaches $0 + M = M$, applied to $P_0 = \bar{B}/x^L$ parcels, for a PES bill of $\beta(\bar{B}/x^L)M$. At N_0^E exactly, no PES is written and the monitoring bill is zero. Both branches share the easement cost $N_0^E C_E$, which is continuous, so the downward jump in total cost is $\beta(\bar{B}/x^L)M > 0$. □

Proof of Proposition 8

The easement-only cost $TC_0 = N_0^E \cdot C_E$ is independent of M. Let $TC^*_{\text{PES}}(M) = \min_{N^E < N_0^E} TC(N^E; M)$ be the optimized PES-inclusive cost. By the envelope theorem, $dTC^*_{\text{PES}}/dM = \beta P^*(M) > 0$, since the indirect effect through the optimal N^E vanishes at the optimum; hence TC^*_{PES} is strictly increasing in M. At $M = 0$, Proposition 4 gives $TC^*_{\text{PES}}(0) < TC_0$. As $M \to \infty$, $TC^*_{\text{PES}}(M) \to \infty$ while

TC_0 is fixed. By continuity and strict monotonicity there is a unique M^* with $TC^*_{\text{PES}}(M^*) = TC_0$; for $M < M^*$ the PES-inclusive strategy is cheaper, and for $M \geq M^*$ easement-only is cheaper. □

Quadratic specialization. With $\gamma(x) = c(x - x^L)^2$, the interior first-order condition (25) gives $x^* = \sqrt{(x^L)^2 + [M + C_E/(\beta\theta)]/c}$, and M^* solves $TC^*_{\text{PES}}(M^*) = N_0^E C_E$ in closed form given the remaining parameters.

C Robustness to Distributional Assumptions

The baseline model specializes to a uniform offer distribution, $d \sim U[0, \bar{D}]$. This appendix derives the key objects—θ, C_E, the conditional development surplus $\mathbb{E}[d - \pi^L \mid d > \pi^L]$, and the first-order condition—under three alternatives: the Beta distribution (which nests the uniform), the log-normal, and the Pareto. The distribution-free first-order condition (16) holds in every case; the distributions differ only in the shape of the comparative statics, not in the core results.

C.1 Beta Distribution on $[0, \bar{D}]$

Let $d = \bar{D} \cdot Z$ with $Z \sim$ Beta(α, β), $\alpha, \beta > 0$. The uniform is the case $\alpha = \beta = 1$; right-skewed offers (most low, few high) correspond to $\alpha < \beta$. With $F(d) = I_{d/\bar{D}}(\alpha, \beta)$ the regularized incomplete beta function,

$$\theta = 1 - I_{\pi^L/\bar{D}}(\alpha, \beta), \tag{34}$$

$$C_E = \bar{D}\,\frac{\alpha}{\alpha+\beta}\big[1 - I_{\pi^L/\bar{D}}(\alpha+1, \beta)\big] - \pi^L\,\theta, \tag{35}$$

$$\mathbb{E}[d - \pi^L \mid d > \pi^L] = \bar{D}\,\frac{\alpha}{\alpha+\beta} \cdot \frac{1 - I_{\pi^L/\bar{D}}(\alpha+1, \beta)}{1 - I_{\pi^L/\bar{D}}(\alpha, \beta)} - \pi^L. \tag{36}$$

The distribution-free first-order condition is unchanged: $\mathbb{E}[d - \pi^L \mid d > \pi^L] = \beta_{\text{disc}}\, h(x^*)$, where β_{disc} is the discount factor (written out to avoid collision with the Beta shape parameter β). The two-channel structure of the development-pressure comparative static survives: a higher $\bar{D}$ raises both θ and the conditional surplus, so the cost and spreading channels both operate and the inverted-U in $N^E(\bar{D})$ can still arise, with its peak shifted by the skewness α/β. The Beta family is attractive because it nests the uniform, retains the $\bar{D}$ parameterization, and accommodates the realistic case in which many parcels face negligible development pressure while a few face intense offers.

C.2 Log-Normal Distribution

Let $\ln d \sim \mathcal{N}(\mu, \sigma^2)$, so d has support $(0, \infty)$ and Φ denotes the standard normal CDF. Then

$$\theta = 1 - \Phi\left(\frac{\ln \pi^L - \mu}{\sigma}\right), \tag{37}$$

$$C_E = e^{\mu+\sigma^2/2}\left[1 - \Phi\left(\frac{\ln \pi^L - \mu - \sigma^2}{\sigma}\right)\right] - \pi^L\left[1 - \Phi\left(\frac{\ln \pi^L - \mu}{\sigma}\right)\right], \tag{38}$$

$$\mathbb{E}[d - \pi^L \mid d > \pi^L] = \frac{e^{\mu+\sigma^2/2}\left[1 - \Phi(\frac{\ln \pi^L - \mu - \sigma^2}{\sigma})\right]}{1 - \Phi(\frac{\ln \pi^L - \mu}{\sigma})} - \pi^L. \tag{39}$$

The distribution-free first-order condition holds. The bounded-support parameterization $\bar{D}$ is replaced by the location and scale (μ, σ): the conditional surplus rises in both, so the cost channel of the development-pressure comparative static is present, and θ falls while the conditional surplus rises in π^L, preserving the ambiguity of $dN^E/d\pi^L$. The log-normal is empirically well motivated, since land values and real-estate prices are approximately log-normal, and its two parameters allow the level and dispersion of development offers to be calibrated separately from transaction data.

C.3 Pareto Distribution

Let $d \sim$ Pareto(π^L, α) with shape $\alpha > 1$ and scale π^L, so $F(d) = 1 - (\pi^L/d)^\alpha$ for $d \geq \pi^L$. Under this specification every offer exceeds working-land profit, so development is certain absent an easement. To allow for parcels facing no development pressure, use a mixture: with probability p_0 the parcel receives no viable offer ($d < \pi^L$); with probability $1 - p_0$ it draws from the Pareto. Then

$$\theta = 1 - p_0, \qquad \mathbb{E}[d - \pi^L \mid d > \pi^L] = \frac{\alpha \pi^L}{\alpha - 1} - \pi^L = \frac{\pi^L}{\alpha - 1}, \qquad C_E = \frac{(1 - p_0)\pi^L}{\alpha - 1}. \tag{40}$$

The distribution-free first-order condition is $\pi^L/(\alpha-1) = \beta\, h(x^*)$. The conditional surplus $\pi^L/(\alpha - 1)$ is increasing in π^L, preserving the ambiguity of $dN^E/d\pi^L$. Development pressure is now parameterized by tail thickness α and offer probability p_0: a lower α thickens the tail and raises the conditional surplus, while a lower p_0 reduces the share of parcels at risk. A higher α (thinner tails) lowers the conditional surplus and unambiguously reduces N^E, because $\theta = 1 - p_0$ does not move with α, so only the cost channel operates through α. The separation of tail thickness from development probability gives the Pareto mixture richer comparative statics than the single-parameter uniform, and its heavy tail captures the empirical pattern in which a few parcels near the urban fringe draw very high offers while most face modest pressure.

C.4 Summary

Table 2 collects the key objects and comparative statics. All four specifications satisfy the distribution-free first-order condition (16), and all core results—PES dominance, spreading, no easement-only without frictions, the monitoring-induced easement-only regime, and the heterogeneity targeting results—hold regardless of the offer distribution, because they depend only on the convexity of γ and the structure of the cost-minimization problem.

Table 2. Results and Comparative Statics Under Alternative Offer Distributions

	General F	**Uniform**	**Beta(α, β)**	**Log-Normal**	**Pareto Mix**
Distribution					
Support	$[0, \bar{D}]$	$[0, \bar{D}]$	$[0, \bar{D}]$	$(0, \infty)$	$\{0\} \cup [\pi^L, \infty)$
Pressure param.	—	$\bar{D}$	$\bar{D}$	μ, σ	α, p_0
Nests uniform?	Yes	—	Yes	No	No
Key objects					
θ	$1 - F(\pi^L)$	$\frac{\bar{D}-\pi^L}{\bar{D}}$	$1 - I_{\pi^L/\bar{D}}(\alpha, \beta)$	$1 - \Phi(\frac{\ln \pi^L - \mu}{\sigma})$	$1 - p_0$
$\mathbb{E}[d - \pi^L \mid d > \pi^L]$	$\frac{C_E}{\theta}$	$\frac{\bar{D}-\pi^L}{2}$	(inc. beta fn.)	(trunc. log-N)	$\frac{\pi^L}{\alpha-1}$
Core results (all distribution-free)					
PES dominates acquisition	✓	✓	✓	✓	✓
PES spreading	✓	✓	✓	✓	✓
No easement-only	✓	✓	✓	✓	✓
Distribution-free FOC	✓	✓	✓	✓	✓
Monitoring $\Rightarrow$ easement-only	✓	✓	✓	✓	✓
Heterogeneity targeting	✓	✓	✓	✓	✓
Comparative statics					
$dN^E/d\bar{B}$	+	+	+	+	+
dN^E/dN	−	−	−	−	−
$dN^E/d(\text{pressure})$	ambig.	ambig.	ambig.	ambig.	ambig.
$dN^E/d\pi^L$	ambig.	ambig.	ambig.	ambig.	ambig.
Inverted-U in pressure	possible	possible	possible	possible	possible

Notes: "Pressure" refers to the distribution-specific development-pressure parameter: $\bar{D}$ for Uniform and Beta, μ for Log-Normal, $1/\alpha$ for Pareto. The distribution-free first-order condition is $\mathbb{E}[d - \pi^L \mid d > \pi^L] = \beta\, h(x^*)$. "No easement-only" refers to the model without monitoring costs. The signs of $dN^E/d\bar{B}$ and dN^E/dN follow from $\partial P/\partial N > 0$ and $\partial x^*/\partial \bar{B} > 0$, which hold for any F.

References

Baillie, J. and Y.-P. Zhang (2018). Space for nature. *Science 361*(6407), 1051–1051.

Byrd, K. B., P. Alvarez, B. Sleeter, L. Flint, A. Flint, R. Cameron, and J. Creque (2019). Scenarios of climate adaptation potential on protected working lands from management of soils. *Environmental Research Letters 14*(10), 104001.

Coase, R. H. (1937). The nature of the firm. *economica 4*(16), 386–405.

Conrad, J. M., C. P. Gomes, W.-J. van Hoeve, A. Sabharwal, and J. F. Suter (2012). Wildlife corridors as a connected subgraph problem. *Journal of Environmental Economics and Management 63*(1), 1–18.

Hart, O. and J. Moore (1990). Property rights and the nature of the firm. *Journal of Political Economy 98*(6), 1119–1158.

Hart, O. D. (1988). Incomplete contracts and the theory of the firm. *JL Econ. & Org. 4*, 119.

Jack, B. K., C. Kousky, and K. R. Sims (2008). Designing payments for ecosystem services: Lessons from previous experience with incentive-based mechanisms. *Proceedings of the National Academy of Sciences 105*(28), 9465–9470.

Kremen, C. and A. Merenlender (2018). Landscapes that work for biodiversity and people. *Science 362*(6412), eaau6020.

Lambin, E. F. and P. Meyfroidt (2011). Global land use change, economic globalization, and the looming land scarcity. *Proceedings of the National Academy of Sciences 108*(9), 3465–3472.

Meyfroidt, P., A. de Bremond, C. M. Ryan, E. Archer, R. Aspinall, A. Chhabra, G. Camara, E. Corbera, R. DeFries, S. Díaz, et al. (2022). Ten facts about land systems for sustainability. *Proceedings of the National Academy of Sciences 119*(7), e2109217118.

Parker, D. P. (2004). Land trusts and the choice to conserve land with full ownership or conservation easements. *Natural Resources Journal*, 483–518.

Parker, D. P. and W. N. Thurman (2018). Tax incentives and the price of conservation. *Journal of the Association of Environmental and Resource Economists 5*(2), 331–369.

Reynolds, M. D., B. L. Sullivan, E. Hallstein, S. Matsumoto, S. Kelling, M. Merrifield, D. Fink, A. Johnston, W. M. Hochachka, N. E. Bruns, et al. (2017). Dynamic conservation for migratory species. *Science Advances 3*(8), e1700707.

Salzman, J., G. Bennett, N. Carroll, A. Goldstein, and M. Jenkins (2018). Payments for ecosystem services: Past, present and future. *Tex. A&M L. Rev. 6*, 199.

Shah, P. and A. W. Ando (2016). Permanent and temporary policy incentives for conservation under stochastic returns from competing land uses. *American Journal of Agricultural Economics 98*(4), 1074–1094.