Dynamic Investment in Ecosystem Restoration

Jonah Danziger* Andrew J. Plantinga†

August 2025

Abstract

There is growing acknowledgment of the need to restore degraded environments. This paper studies optimal investment in ecosystem restoration under environmental change. We develop an optimal control model of the restoration decision to explicitly characterize the optimal extent and timing of restoration given time-dependent marginal damages. We provide the first results on optimal dynamic investment in ecosystem restoration, highlighting the important role that growth in restored patches plays in shaping the time profile of investment. We then apply the model in a numerical simulation of coastal wetlands restoration in Huntington Beach, California, that accounts for projected sea-level rise, uncertainty over flooding severity, and the option to abandon damage properties. Our results show that early investments in restoration are optimal in order build up a wetlands stock that can mitigate future flooding damages exacerbated by sea-level rise. We find large option values associated with delaying irreversible decisions to abandon damaged properties.

*Department of Economics, University of California, Santa Barbara, email: jonahdanziger@ucsb.edu

†Bren School of Environmental Science and Management, University of California, Santa Barbara, email: plantinga@bren.ucsb.edu

The authors acknowledge funding from the University of California, Office of the President, National Laboratory Fees Research Program, and helpful comments from Alexander Abajian, Chris Costello, Olivier Deschenes, Robert Heilmayr, Charlie Kolstad, Mark Lubell, Kyle Meng, Antony Millner, and David Zilberman, and participants in the 2024 World Conference on Natural Resource Modeling.

"The next century will, I believe, be the era of restoration in ecology." – E.O. Wilson (1999)

1 Introduction

Healthy ecosystems are the foundation for the flow of services that sustain life on Earth. Whereas early conservation efforts focused on preserving unique environments, there is growing recognition of the need to restore damaged and degraded ecosystems (Blignaut et al., 2014). Ecosystem restoration is the process of assisting the recovery of environments that have been deleteriously altered by human activities. It can be a passive or active endeavor. Passive restoration is when a degraded ecosystem is set aside and allowed to recover on its own. Marine protected areas are one prominent example. Active restoration involves deliberate human interventions, such as planting of native trees, grasses, and corals or rewetting of peatlands. It also includes changes in management practices, such as crop rotations to restore agricultural soils and controls on the timing and intensity of livestock grazing. Whether passive or active, the goal is to restore ecosystem functionality and service flows.[1] To this end, there are currently initiatives throughout the world promoting ecosystem restoration (Mirzabaev and Wuepper, 2023), including at the global scale the United Nations Decade on Ecosystem Restoration.[2]

Economic studies of ecosystem restoration have emphasized estimation of the benefits and costs of restoration, *ex ante* frameworks for evaluating projects, and the design of policies to promote restoration efforts (Mirzabaev and Wuepper, 2023). Missing from the literature is a complete analysis of optimal ecosystem restoration. The manager faces a fundamentally dynamic problem, deciding when and how much restoration to undertake in order to influence the evolving state of the ecosystem. In a world with environmental (e.g., climate) change, the system is often characterized by stochasticity and non-stationarity. For example, the decision to restore coastal wetlands for flood protection must account for the growth over time in wetland plants, the accumulation of sediments, and uncertain threats from sea-level rise. Although optimal ecosystem restoration has been analyzed in static models (e.g., Neeson et al. (2015)), to our knowledge, no studies examine the dynamics of optimal restoration.

The central research question addressed in this study is: what should be the extent and timing of ecosystem restoration under environmental change? High upfront costs and increasing marginal benefits of restoration argue for delaying projects, but if a restored patch grows over time, and thus provides higher benefits the earlier it is restored, near-term investment becomes more valuable. We develop an optimal

[1]Some authors make a distinction between ecosystem restoration and rehabilitation (e.g., Arneth et al. (2021)), where the former refers to actions that support ecological processes and the latter to actions that enhance ecosystem service flows. Our analysis applies to restoration and rehabilitation activities and so the distinction is not important. As such, we follow Blignaut et al. (2014) and use the single term "restoration."

[2]https://www.decadeonrestoration.org/

control model to characterize optimal restoration decisions through time. The decision-maker chooses the area of patches to restore in each period to maximize the flow of benefits from the ecosystem net of losses from unrestored areas and restoration costs. In a standard natural resource model, the state variable (e.g., the biomass stock) is a function only of time. In our problem, we account for the age of each restored patch as this determines the level of benefits it provides and the rate at which it grows. An age-class model is needed to represent many types of ecosystem restoration. The accumulation of nutrients in degraded farmland soils, the service flows from reforestation, and the re-establishment of native grassland plants depend on which patches have been restored and patch-specific times since restoration.[3] In these examples, the change in the state of the system depends on the evolution of individual restored patches, rather than an aggregate state variable such as total restored area. In our model, the state variable is a function of time and age, and the state equation is a partial differential equation. Adapting a model of vintage capital, discussed below, we are able to obtain an analytical solution to the optimal restoration problem. In each period, the marginal cost of restoration equals the shadow value of an age 0 restored patch, equal to the present value sum of avoided damage over the lifetime of the patch, accounting for its growth over time and the evolving marginal damages from environmental change.

Our theory of optimal ecosystem restoration is related to models developed in a number of previous papers in environmental and natural resource economics. Smith et al. (2009) study the optimal restoration of beaches with imported sand (i.e., beach nourishment). Beaches erode over time at an exogenous rate, which is analogous to the growth of a restored patch in our model. However, nourishment is assumed to always restore a beach to a fixed volume of sand, which means that only the timing of nourishment events, and not the amount of nourishment, is chosen by the decision-maker. Thus, the problem is similar to finding the optimal timber rotation with standing stock benefits (e.g., Hartman (2018)). Our model has some common features with models of optimal pollution clean-up (e.g., Caputo and Wilen (1995)). In Lappi (2018), the decision-maker determines the optimal clean-up of polluted sites to minimizes damages plus costs, where pollution decays at an exogenous rate. The assumption that benefits and costs are additive across sites, each site is treated only once, and each site must be fully remediated means that only the timing of clean-up at each site must be chosen. Our model is more general in that the decision-maker chooses the amount of restoration, marginal benefits depend on the cumulative amount of restoration, and growth of restored patches is age-dependent.

We draw inspiration from models of vintage capital (Benhabib and Rustichini, 1991; Goetz et al., 2008). In particular, Feichtinger et al. (2006) solve the complete dynamic optimization problem for a firm who buys or sells machines of different ages to maximize discounted profits. The state variable in their model

[3] Additional examples are found with the restoration of coral reefs, wetlands, and peatlands.

is $K(t, a)$, the stock of capital in time t of age a. The state equation is then given by the linear partial differential equation:

$$\frac{\partial K(t,a)}{\partial t} + \frac{\partial K(t,a)}{\partial a} = I(t,a) - \delta(a)K(t,a) \tag{1}$$

where $I(t, a)$ is the number of machines of age a bought or sold in t and $\delta(a)$ is the age-specific depreciation rate. If capital is homogeneous, (1) simplifies to the familiar expression for capital stock dynamics: $\dot{K}(t) = I(t) - \delta K(t)$. Feichtinger et al. (2006) show that the model with state dynamics given by (1) can be solved with the addition of a boundary Hamiltonian, which represents the flow value of new (age 0) machines, and application of the Method of Characteristics to convert the adjoint equation to an ordinary differential equation. The authors obtain an explicit expression for optimal investment rate through time.[4] We adapt the Feichtinger et al. (2006) model to the problem of optimal ecosystem restoration, where our state equation depends on time as well as the age of restored patches.

To further explore the insights from our theoretical analysis, we develop a numerical model applied to coastal wetlands restoration. Coastal communities throughout the world are increasingly vulnerable to flooding as the result of sea-level rise (Kulp and Strauss, 2019). Although hardened structures such as seawalls have traditionally been used to manage these risks, there is increasing interest in the use of nature-based solutions to reduce wave energy and storm surge (Sun and Carson, 2020; Fairchild et al., 2021). The dynamic simulation relaxes two assumptions in the optimal control model. First, we model damages as stochastic rather than deterministic, allowing us to explore the implications of uncertainty for optimal ecosystem restoration. Second, we consider the possibility that an area may be optimally abandoned when the costs of restoration are too high relative to net benefits from the ecosystem. We model the abandonment decision as irreversible, which when combined with uncertain damages, gives rise to option value — namely, the decision to restore an area preserves the option to restore or abandon it in the future. In addition to coastal wetlands restoration, option value is likely to arise in other settings where large shocks (e.g., severe wildfire) makes recovery of the system extremely expensive (see, e.g., Syphard et al. (2022)).

We apply the numerical model to Huntington Beach, California. Like many parts of coastal California, much of the Huntington Beach area was historically in wetlands (Stein et al., 2014; Grossinger et al., 2011) and is now vulnerable to sea-level rise. As shown in Figure 1, significant amounts of land currently in developed uses are projected to be under water by the end of the century under a 5-foot sea-level rise scenario. Our simulation model accounts for the morphology of the coastal region, wave characteristics, sea-level rise projections, and the value of assets at risk. Uncertainty in flooding damages arises from stochastic variation in wave heights, which we measure using buoy data from the National Oceanic and Atmospheric Admin-

[4]Xepapadeas and de Zeeuw (1999) analyze a model with a similar structure to understand how environmental regulation affects the composition of a firm's capital stock. However, they characterize only the steady-state optimality conditions.

istration (NOAA). Our numerical simulation solves for the optimal restoration of wetlands, which protect developed areas from storm surges and sea-level rise. When developed areas are damaged by flooding, the decision-maker must decide whether to repair properties or abandon them. We treat the abandon decision as irreversible, which gives rise to an option value associated with maintenance of properties. Simulation results highlight the important role that wetlands growth plays in shaping the dynamics of optimal restoration. Even though flooding risks are relatively low in the near term, costly investments in restoration are undertaken so that growth in wetlands will eventually protect properties from sea-level rise. The possibility of abandoning damaged properties means that severe flooding events can actually reduce the value and optimal amount of restoration investments. For properties in our study region, option values associated with delaying irreversible abandonment decisions are found to be worth a total of $1.6 billion.

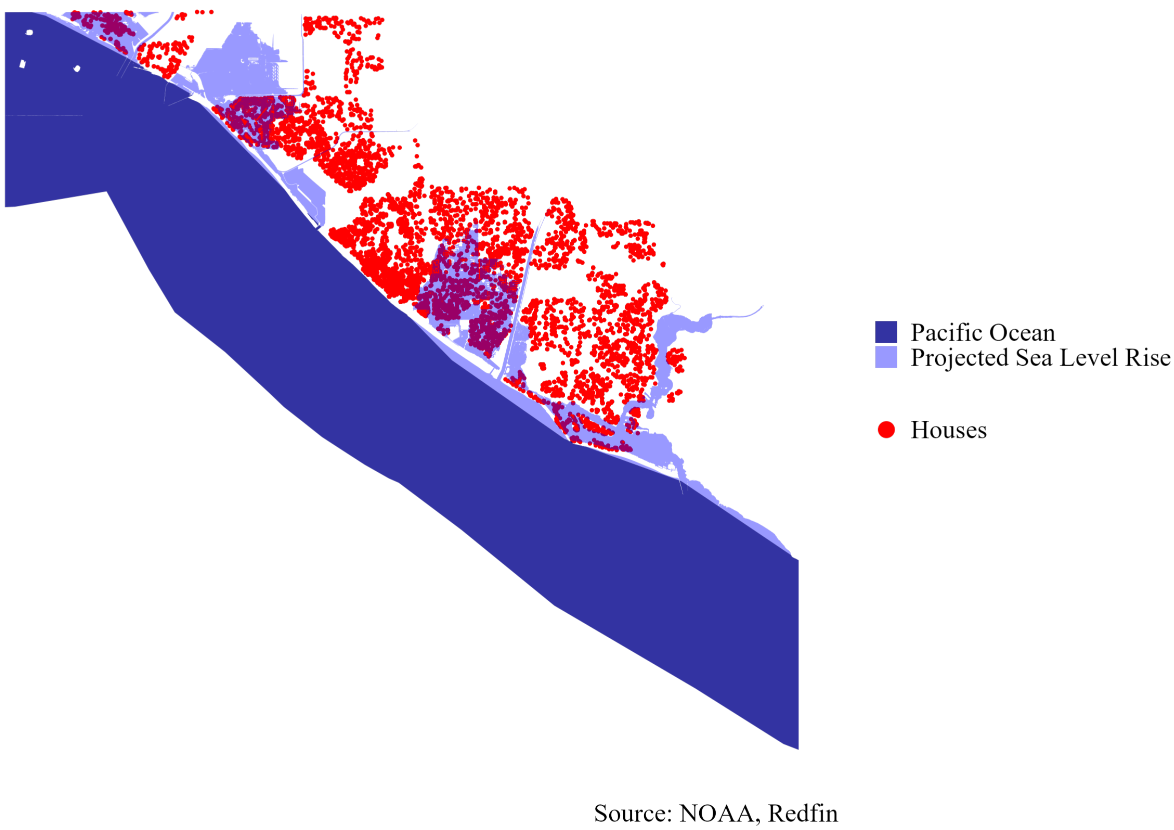

Figure 1: Huntington Beach, California, Under a Projected 5-foot Sea-level Rise

The next section presents an optimal control model of ecosystem restoration. In addition to deriving an explicit solution, we provide numerical results to elucidate how the time path of the control variable changes under different conditions. Section 3 presents the empirical application to wetlands restoration in

Huntington Beach, CA, under sea-level rise. Section 4 discusses the optimal solution and presents sensitivity analyses and option value results. A final section provides discussion and conclusions.

2 The Resource Manager's Problem

We model the decision problem for a resource manager who seeks to maximize the present value of net benefits from ecosystem restoration under environmental change.[5] The control problem is given by:

$$\begin{aligned} \max_{y(t)} \int_0^T e^{-rt}\big[\int_0^t b(t) - p(t)(\bar{x} - x(t,\alpha))d\alpha - C(y(t))\big]dt \quad \text{s.t.} \\ \frac{\partial x(t,\alpha)}{\partial \alpha} + \frac{\partial x(t,\alpha)}{\partial t} = x(t,\alpha)f(\alpha), \quad \alpha \in [0,t], \quad x(t,0) = y(t), \quad \lambda(T,\alpha) = 0 \end{aligned} \tag{2}$$

The benefit flow in time t from a fully restored ecosystem is given by $b(t)$. This amount is reduced by damages from unrestored areas, equal to the product of time-dependent marginal damages, $p(t)$, and the gap between the current state of the system, $x(t,\alpha)$, and a fully-restored system, $\bar{x}$. The state variable in t, $x(t,\alpha)$, depends on the age α of the patch, where the inner integral in (2) sums over ages 0 to t.[6] Marginal damage changes over time to capture effects of environmental change, such as rising sea levels or more favorable conditions for invasive species. In the outer integral, from time 0 to the end of the planning horizon T, the cost of restoration $C(y(t))$ is a convex function that increases with the amount of restoration in t, $y(t)$. The state equation is the rate of change in the restored stock of a given age, which is a function solely of the growth in the existing stock, where $f(\alpha)$ is a concave function. We assume that the manager in our model can invest in new restoration but cannot augment existing restored areas.[7] Thus, the amount of restoration in time t, $y(t)$, equals the stock of an age 0 restored patch in time t, $x(t,0)$. Finally, the transversality condition requires that the shadow value of the stock in the final period T, $\lambda(T,\alpha)$, is equal to zero.

To solve the model, we need to specify current-value and boundary Hamiltonians. The current-value Hamiltonian equals the flow of value from a restored patch of age α at time t:

$$\tilde{H}(x,y,\lambda) = b(t) - p(t)(\bar{x} - x(t,\alpha)) + \lambda(\alpha,t)f(\alpha)x(\alpha,t) \tag{3}$$

[5]We assume that the resource manager has selected a particular restoration project for evaluation. A separate question is whether to pursue this project, an alternative project, or any project at all. In Section 5, we discuss how our analysis is needed to address this broader issue.

[6]Note that $\bar{x}$ does not affect the solution to the optimization problem. As such, we are not bound by an assumption that damages are completely eliminated when $x(t,\alpha) = \bar{x}$. In addition, we could solve the integral $\int_0^t x(t,\alpha))d\alpha$ and replace it with $X(t)$. However, as shown in Appendix A.1, we can write the state equation as a function of $X(t)$ only for a special case of $f(\alpha)$ and so we keep damages in the form in (2).

[7]This is in contrast to the Feichtinger et al. (2006) model (see their equation 1).

The boundary Hamiltonian captures the flow of value of an age 0 patch:

$$H_0 = -C(y(t)) + \lambda(t,0)y(t) \tag{4}$$

The shadow value, $\lambda(t,\alpha)$, is the present value of restoring a marginal amount of the ecosystem. Its dependence on age is a generalization of standard models of natural resources in which the shadow value depends only on time. The maximum conditions and adjoint equation for the model are given by:

$$\frac{\partial \tilde{H}}{\partial y} = 0 \tag{5}$$

$$\frac{\partial H_0}{\partial y} = 0 = -C'(y(t)) + \lambda(t,0) \tag{6}$$

$$\frac{\partial \lambda}{\partial t} + \frac{\partial \lambda}{\partial \alpha} = r\lambda - \frac{\partial \tilde{H}}{\partial x} = -p(t) + \lambda(r - f(\alpha)) \tag{7}$$

We use equations (6) and (7) to solve for $y(t)$, $x(t,\alpha)$, and $\lambda(t,\alpha)$. The Method of Characteristics is used to transform the partial differential equation in (7) into an ordinary differential equation (ODE). Once in the form of an ODE, we can use an integrating equation and integration by parts to obtain an explicit solution to the control problem in (2):

$$\lambda(t,\alpha) = \int_t^T e^{-\int_t^s (r - f(\alpha+\rho-t))d\rho} p(s)ds \tag{8}$$

$$C'(y(t)) = \lambda(t,0) \tag{9}$$

$$x(t,\alpha) = x(t-\alpha,0)e^{\int_0^\alpha f(\rho)d\rho} \tag{10}$$

All derivations are found in see Appendix A.2. The shadow value in (8), the value of a marginal increase in a patch of age α in time t, is equal to the present value sum of the damage prevented over the lifetime of the restored patch, accounting for its growth over time and the evolving marginal damages from environmental change. It then follows from the maximum condition in (6) that restoration should continue to the point at which the marginal cost of restoration equals the marginal benefit of establishing an age 0 patch (equation (9). Lastly, the stock of a patch of age α at time t is equal to the amount that was initially restored, augmented with α years of growth at different growth rates given by $f(\alpha)$ (equation (10).

If we adopt specific functional forms, we can gain additional insights into the timing and magnitude of optimal investment in restoration. We set our cost function to be a quadratic function over the amount restored (see Appendix A.3 for details). Growth is an inverse function of the age of the restored patch and damages are a logistic function of time. With these functional forms, we can derive the key result in equation 9:

$$2(y(t)-1)=\int_0^T (s-t+1)e^{-r(s-t)}\frac{N}{1-e^{-B(s-A)}}ds \tag{11}$$

where T, r, N, B, and A are parameters. With this result, we can adjust the parameters to gain insight into the optimal choice of restoration under different conditions. We first establish a baseline level of restoration, and then consider the implications of higher damages and damages that occur further into the future (Figure 2). With the baseline parameters, restoration is highest in the early years and tapers off as the end of the time horizon ($T = 100$) approaches. Even though damages are low initially, a high level of initial investment takes advantage of growth in the stocks over time. If damages are shifted to further in the future, then restoration effort increases initially before declining. In the initial period, the discounted future damages are lower than in the base case, causing a delay in optimal investment. The shift up in the level of restoration in every period due to higher damages is intuitive as is the more rapid decline in investment since the manager will inevitably make zero investment in the last period. Finally, a higher interest rate lowers the present value of future damages in every period and, thus, encourages not only lower initial amounts of restoration, but a flatter overall restoration schedule.

The analytical model in (2) departs from standard natural resource models by allowing for age-dependent growth. Thus, a natural question is how this generalization affects the optimal solution. We show in Appendix A.1 that when growth is constant (i.e., $f(\alpha) = \theta$), the state equation with age classes can be rewritten as a function of aggregate restoration $X(t) = \int_0^t x(t,\alpha)d\alpha$, as in a standard model. In Figure 3, we compare optimal investment in restoration under age-dependent growth to investment with constant growth.[8] In the latter case, we fix $f(\alpha) = \frac{1}{1+\alpha}$ at a range of values $f(13), f(14)$, etc. The solution for the base case with age-dependent growth in Figure 2 is reproduced and shown in red. Under the assumption of constant growth, the optimal restoration path becomes convex. Compared to the base case, there is no incentive to restore new patches to take advantage of the higher growth rates that occur at younger ages.[9] Thus, investment occurs early and diminishes rapidly over time rather than being spread out over the entire time horizon. When growth is slow (larger values of α), there is little investment in restoration.

[8] The same functional form assumptions and parameter values used to produce Figure 2 are used here; see Appendix A.3 for details.

[9] By assuming constant growth, $f(\alpha) = \theta$, our model has the key feature of a standard renewable resource model; i.e., an aggregate state variable $X(t)$ that is not differentiated by age classes. However, in our model the assumption of constant growth implies that the growth rate in the aggregate stock, $\dot{X}/X$, is constant, whereas in the standard model the growth rate in the stock typically declines as X approaches the carrying capacity (see Appendix A.1 for details). We would expect a declining growth rate to further reduce the incentive to establish new patches and flatten the curves shown in Figure (2).

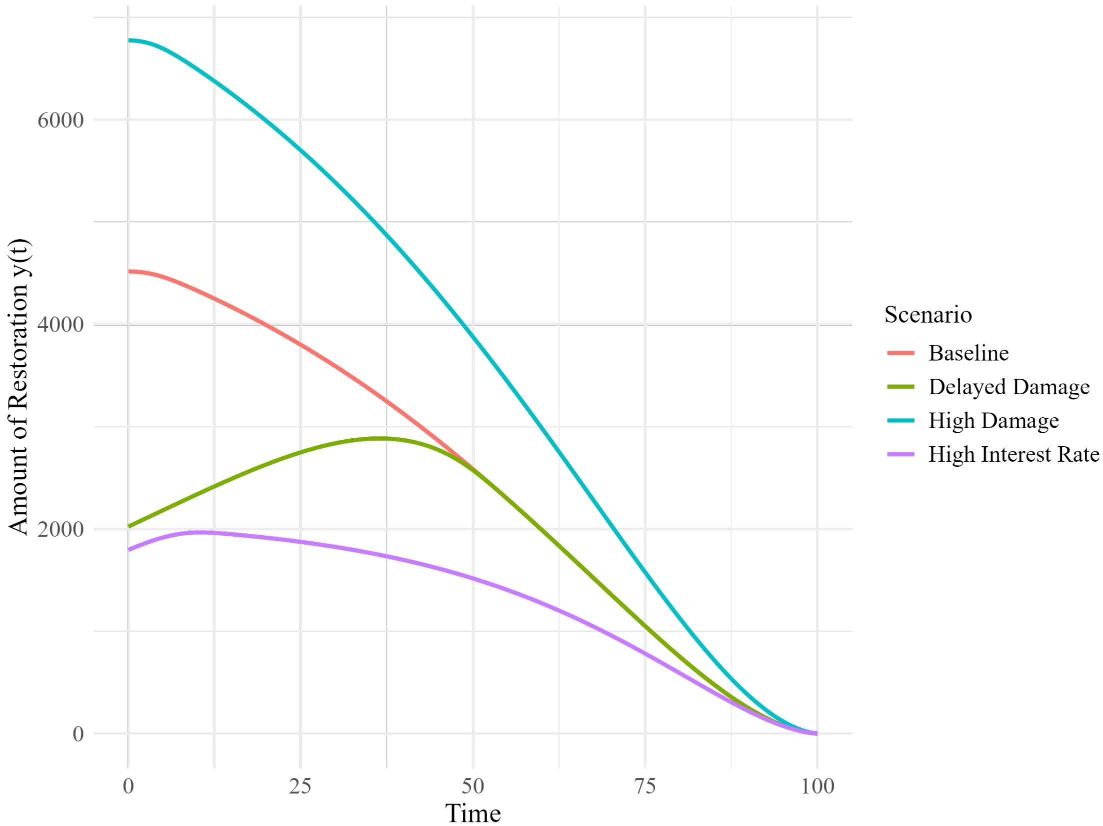

Figure 2: Analytical Solution with Imposed Functional Forms

3 Empirical Application

Approximately 10% of the world's population live in low-elevation coastal zones susceptible to flooding (Reimann et al., 2023). Coastal habitats provide a natural defense against flooding, while avoiding negative environmental effects of hardened structures and providing a range of ecosystem services (Arkema et al., 2013). Coastal flooding often occurs during storms, when water from the ocean is pushed inland in a phenomenon known as storm surge. Sea-level rise exacerbates flooding damages by increasing the inland advance of storm surges. Coastal wetlands help to mitigate storm surge by reducing wave energy and run up (Narayan et al., 2016). We develop an empirical simulation to maximize the present value of expected net benefits from wetlands restoration in Huntington Beach, California. Like many other coastal regions in the United States, this area once had extensive wetlands and tributaries that were filled in and developed following Euro-American settlement. Thus, today the areas that were historically in wetlands tend to be low-lying and vulnerable to the combined effects of sea-level rise and storm surge.

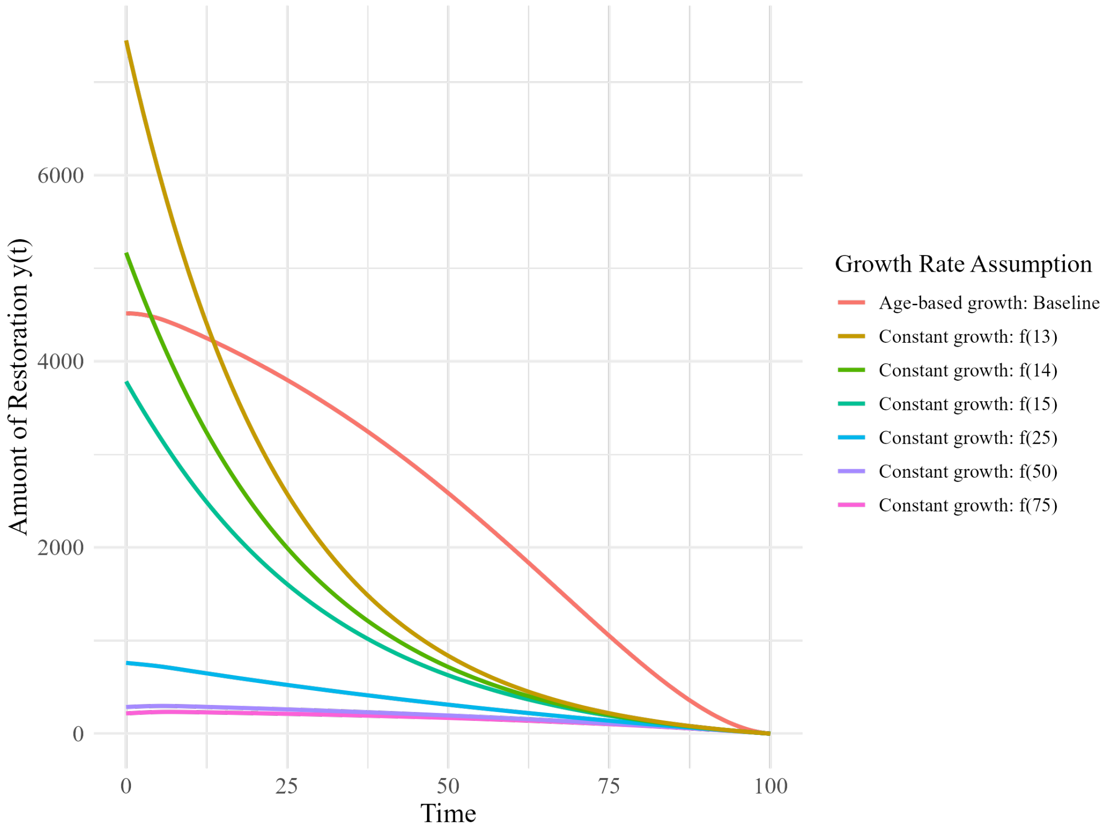

Figure 3: Restoration Under Alternative Growth Assumptions

3.1 Empirical Model

Similar to the analytical model in section (2), the resource manager chooses the amount of wetlands restoration in each period to maximize the present discounted value of coastal real estate net of restoration costs. Define $\mathbf{y}_0 = \{y_0, y_1, ..., y_{T-1}\}$ as the set of restored areas over time. At time $t+1$, the stock of restored areas is:

$$X_{t+1} = \sum_{s=0}^{t} y_s \prod_{\alpha=0}^{t-s} (1 + f(\alpha)) + X_0 \tag{12}$$

where $f(\alpha)$ is the growth in an age α patch of unit area and X_0 represents the existing stock of wetlands at time 0. The product term in (12) is the discrete-time analog to the exponential term in (10).

There are $i = 1, ..., N$ properties that generate rents r_i in each year. Properties are subject to iid flooding shocks z_{it} that cause damages $D(z_{it}, X_t)$, where damages are decreasing in the aggregate stock of restored wetlands: i.e., $D_1 > 0, D_2 < 0$. The resource manager decides whether to abandon ($d_{it} = 0$) or maintain ($d_{it} = 1$) each property in each period. Define a_{it} as a state variable indicating whether property i has been abandoned ($a_{it} = 0$) or maintained ($a_{it} = 1$) as of the start of period t. a_{it} evolves according to:

$$a_{it+1} = a_{it} d_{it} \tag{13}$$

Note that once a property is abandoned, it remains in that state forever (i.e., abandonment is irreversible).

The timing of the model is as follows: 1) period t begins with states X_t and a_{it} given, 2) restoration y_t is chosen, 3) rents are collected for each property i if $a_{it} = 1$, 3) the shocks z_{it} happen, 4) decisions to abandon or maintain each property i, d_{it}, are made, and 5) period $t+1$ begins with new values of X_{t+1} and a_{it+1} according to (12) and (13). The decision to maintain property i means that costs $D(z_{it}, X_t)$ are incurred (i.e., the damages are reversed). If abandoned, these costs are avoided. Thus, when maintained, the net rent from property i in t is:

$$\pi_{it} = a_{it}[r_i - d_{it} D(z_{it}, X_t)] \tag{14}$$

If abandoned in t, the rent is $\pi_{it} = a_{it} r_i$.

Conditional on a restoration plan $\mathbf{y}_t = \{y_t, y_{t+1}, ..., y_{T-1}\}$, which determines the stock of restored area over time, $\mathbf{X}_{t+1} = \{X_{t+1}, X_{t+2}, ..., X_T\}$, the solution to the abandon/maintain subproblem for property i is given by Bellman's equation:

$$V(X_s, a_{is}, z_{is}) = \max_{d_{is}} \pi_{is} + \delta E_z V(X_{s+1}, a_{is+1}, z_{is+1}) \tag{15}$$

for $s = t, t+1, ..., T-1$, X_t and a_{it} given, and subject to (13) and (14). δ denotes the discount factor. In period T, each property yields the salvage value $S(X_T, a_{iT}) = \frac{a_{iT} E_z[r_i - D(z_{iT}, X_T)]}{1-\delta}$. If a property has been maintained until T (i.e., $a_{iT} = 1$), $S(X_T, a_{iT})$ equals the present discounted value of an infinite sum of expected rents minus damages in period T. In (15), $V(X_t, a_{it}, z_{it})$ is the present discounted value of expected net rents from property i given optimal abandon and maintain decisions. With this function, we can write the optimal restoration problem in time 0 as:

$$\max_{\mathbf{y}_0} \sum_{i=1}^{N} E_z[V(X_0, a_{i0}, z_{i0})] - \sum_{t=0}^{T-1} \delta^t C(y_t) \tag{16}$$

subject to (12), $a_{i0} = 1 \ \forall i$, and X_0 given, and where $C(y_t)$ is the cost of restoring an area of y_t. The expectation is needed for $V(\cdot)$ because y_t is chosen before the z_{it} are observed.

The solution to the problem in (16) is $\mathbf{y}_0^*(\mathbf{a}_0, X_0) = \{y_0^*, y_1^*, ..., y_{T-1}^*\}$ where $\mathbf{a}_0 = \{a_{10}, a_{20}, ..., a_{N0}\}$ and X_0 are the initial values of the state variables. In words, the solution is the optimal set of restored areas from time 0 to $T-1$ given the initial states in $t = 0$. Because the y_t are chosen before the z_{it} occur, the shocks are "integrated out" via the expectation in (16). Therefore, the solution is the optimal choice y_0^* and a set of $T-1$ values $\{y_1^*, y_2^*, ..., y_{T-1}^*\}$ that are optimal in expectation from the perspective of time 0. We show in Appendix A.4 that this solution is equivalent to the solution obtained with backwards induction. We refer to $\mathbf{y}_0^*$ as the optimal *planned* restoration since, with the exception of y_0^*, it is what the manager expects to do given the information available in time 0.

The resource manager implements y_0^*, which determines X_1 by (12). Before period 0 ends, the flooding shocks $\mathbf{z}_0 = \{z_{10}, z_{20}, ..., z_{N0}\}$ occur. Now the manager makes optimal time 0 decisions to abandon or maintain properties ($d_{i0}^* \ \forall i$) based on realizations of $\mathbf{z}_0$ rather than in expectation. By (13), this results in a new set of state variables $\mathbf{a}_1$. Thus, the manager enters period 1 with new "initial conditions" X_1 and $\mathbf{a}_1$, which means the problem in (16) must be resolved to find $\mathbf{y}_1^*(\mathbf{a}_1, X_1) = \{y_1^*, y_2^*, ..., y_{T-1}^*\}$, where y_1^* is the optimal restored area in time 1 and $\{y_2^*, y_3^*, ..., y_{T-1}^*\}$ are optimal in expectation. The manager implements y_1^* and the process continues until the end of the planning horizon at time T.

3.2 Data and Set-up

To obtain a solution to the problem in (16), we define specific functions for flooding damages, net rents, wetlands growth, and restoration costs. We have a sample of 7,509 properties located in Huntington Beach, California. The mean house price, structure value, and monthly rent are \$1.58 million, \$449 thousand, and \$4536, respectively. The median (mean) elevation of the properties is 6.27m (10.17m). As shown in Figure 1, many of these properties are at risk from projected sea-level rise and storm surge.

Flooding damages depend on flooding shocks z_{it} and the stock of restored wetlands X_t. To specify the shock distribution, we first measure wave run-up, or the maximum onshore elevation reached by waves during storm surge, measured relative to the sea level (or still water level). Erikson et al. (2018) provide the empirical formula relating deep-water wave conditions to wave run-up (R) for southern California near-shore topography:

$$R = 1.1\Big(0.35\beta_f(SWH * L_0)^{0.5} + 0.5 * [SWH * L_0(0.563\beta_f^2 + 0.004)]^{0.5}\Big) \tag{17}$$

where β_f is the beach slope, SWH is significant wave height, and L_0 is deep-water wave length. β_f is estimated as the median slope at the shoreline using a grid of points along the coast at Huntington Beach. SWH and L_0 are derived from NOAA buoy data (specifically Station 46222, located offshore from Huntington Beach) for the period 2004 to 2024. SWH is the maximum daily wave height and L_0 is obtained from the formula $L_0 = gT_p^2/2\pi$ where g is acceleration due to gravity and T_p is the peak daily wave period. We use (17) to find the highest daily run-up in each year over the period 2004-2024 and define $\tilde{R}$ as a random variable drawn from this set of maximum run-up values with equal probability.[10] The flooding shock is given by:

$$z_{it} = \tilde{R} + sl_t - elv_i \tag{18}$$

where sl_t is the still water level in year t and elv_i is the elevation of property i. Property i experiences flooding in year t if $z_{it} > 0$. We measure sl_t using a quadratic function fit to the intermediate-high sea-level rise prediction for Huntington Beach available from the NOAA Sea Level Rise Viewer.[11] The elevation of property i is obtain from Redfin.

Flooding damages in year t for property i are specified:

$$\begin{aligned} D(z_{it}, X_t) &= \frac{A}{1 + e^{-B\left[z_{it} - (X_t/\bar{X})\tilde{R}\right] + C}} strval_i \\ &= \frac{A}{1 + e^{-B\left[\tilde{R}(1 - X_t/\bar{X}) + sl_t - elv_i\right] + C}} strval_i \end{aligned} \tag{19}$$

A, B, and C are parameters obtained by a fitting a logistic function to depth-damage data in Davis and Skaggs (1992).[12] The logistic function in (19) measures the share of the structure value, $strval_i$, destroyed by the flood. Larger flooding shocks increase damages (note that $-B < 0$), whereas wetlands restoration decreases damages via the term $(X_t/\bar{X})\tilde{R}$ where $\bar{X}$ represents the stock of a fully-restored wetland. When $X_t = \bar{X}$, the effect of storm surge is completely mitigated and damages only occur as the result of sea-level rise. To measure the structure value, we start with the sale price of properties available from Redfin, which include both the land and structure value. We then use data from Davis et al. (2021) on census tract level land values to back out the structure value of each property.

In specifying the damage function in (19), we make a simplifying assumption about wetlands restoration. Damages are reduced by the aggregate stock of wetlands, X_t, which are assumed to be placed along the shoreline, but not to take up space currently occupied by houses. A more complicated analysis would consider the trade-off between space for housing and wetlands. We abstract from these spatial complexities to keep the focus on our main interest, the dynamic trade-offs inherent in ecosystem restoration.

We used Redfin data from the past two years of sales and web-scraped estimated rental values to measure r_i, the annual rental income from property i. It is important to use rents instead of home prices because the latter may capitalize future effects of flooding. For a small share of the properties, rental values are unavailable. In these cases, we predict annual rents using a simple hedonic model that specifies rents as a function of house characteristics, distance from the coast, and local demographics (details and results are

[10] We draw only one shock each year and, thus, we assume that damages depend on the most severe annual floods that have occurred historically.

[11] The fitted equation is $sl_t = 0.182 + 0.00592t + 0.000158t^2$ and the data is found at: https://coast.noaa.gov/slr/#.

[12] The estimated values are $A = 0.369$, $B = 0.9497$, and $C = 2.803$.

found in Appendix A.5). The wetlands growth function is specified $f(\alpha) = e^{-\tau\alpha}$, where $\tau = 0.08$ is chosen based on age-dependent growth functions from the ecology literature (Paine et al., 2012). Costs are given by $C(y_t) = \gamma y_t^2$ where, for the base specification, $\gamma = 1$. There is little data on restoration costs, especially for large-scale restoration projects, and so we adopt a quadratic function for $C(y_t)$ that, consistent with the analytical model, has the necessary convexity property. Finally, we set $\delta = 0.97$ and $T = 100$.

3.3 Solution method

To solve the problem in (16), we use gradient methods to search for the optimal restoration vector $\mathbf{y}_0$ that maximizes the expected value of assets net of costs. At each iteration of the search, we use backward recursion to solve the stochastic dynamic programming (SDP) problem in (15) conditional on the candidate set of $\mathbf{y}_0$ values. This nesting of the SDP subproblem within an outer search for $\mathbf{y}_0$ mirrors the method developed by Rust (1987) to estimate dynamic structural models.

We use backpropagation to compute gradients of the objective function with respect to $\mathbf{y}_0$. Backpropagation is a machine learning method commonly used in macroeconomics (Swanson and White, 1997) and non-parametric regression analysis (Cattaneo et al., 2024). It systematically applies the chain rule to efficiently compute derivatives of the objective function. The gradients are then used to adaptively update the choice variables, allowing for stable convergence even in the presence of uncertainty. The advantage of backpropagation over other gradient approximation techniques is that it computes exact gradients efficiently without numerical error. We use gradient methods in combination with dynamic programming due to the computational costs of traditional approaches such as value or policy function iteration. The curse of dimensionality is particularly severe in our setting because a separate state variable is needed for the wetlands volume in each age cohort when growth is nonlinear (see Appendix A.1). Our approach circumvents this problem by choosing the full set of $\mathbf{y}_0$ values in the outer search, allowing us to calculate the aggregate stock of restored areas in each period outside of the SDP algorithm using (12).

4 Numerical Simulation Results

We begin our discussion of results by showing the optimal planned restoration in period 0 — specifically, the values of $\mathbf{y}_0$ that solve (16). The red line in Figure 4a displays the 100-year restoration plan with sea-level rise included. Even though flooding events are becoming more severe over time, investments in restoration decline. Although the finite time horizon plays some role here — namely, restoration in the final period goes to zero because it provides no benefits beyond T — less investment is needed later in the planning period because of growth in the wetlands established earlier. These results underscore the important ways in which

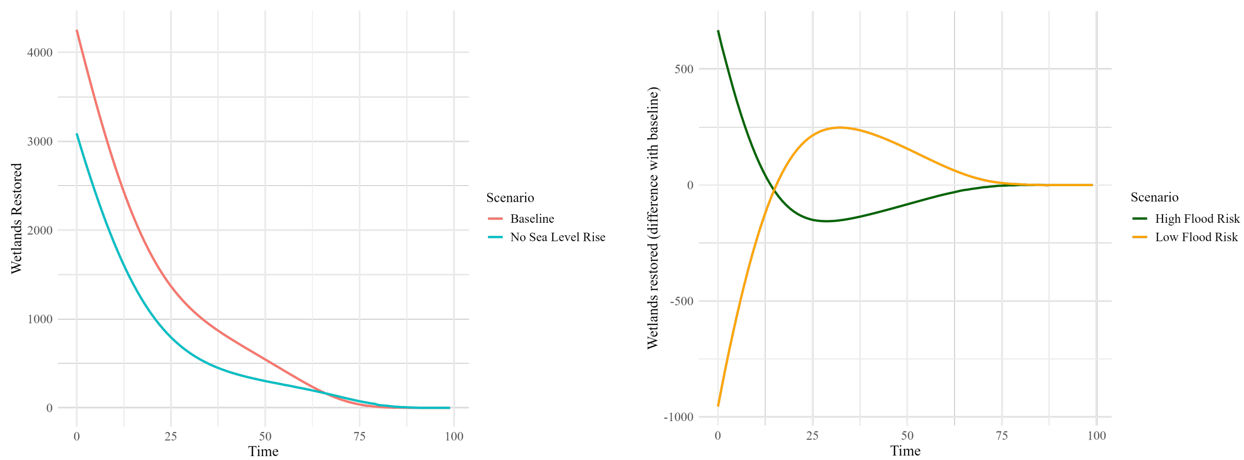

(a) Restoration with and without sea-level rise

(b) High and low storm surge scenarios

Figure 4: Optimal Wetland Restoration Under Baseline and Alternative Scenarios

non-stationary effects of climate change shape optimal restoration strategies. For comparison, the blue line in the figure removes the effects of sea-level rise. Sea-level rise raises investments in restoration because it magnifies the damages from storm surges, but it also encourages investment later in the planning horizon. The flattening of the red curve in about year 40 corresponds to emerging flooding risks as sea levels rise. Investments in restoration are maintained at a relatively high level in order to mitigate expected effects of storm surges intensified by climate change.

Storm intensity has surprising effects on the dynamics of restoration investments. In Figure 4b, we present two scenarios in which the flooding shock distribution is restricted to the five highest and five lowest realizations of $\tilde{R}$. For each scenario, we plot the difference in wetlands restoration relative to the baseline scenario (the red curve in Figure 4a). Although one might expect greater storm intensity to unambiguously raise investment, this is only true early in the planning horizon. In the high flooding risk scenario, initial restoration investment exceeds the baseline level up to about 15 years (see green curve). Early investment ensures a large stock in the future as the result of wetlands growth, providing protection against more extreme storm events. But once these early investments have been made, the marginal value of additional restoration is lower. In addition, greater storm risks mean that more properties are optimally abandoned, reducing the value of restoration. The low flooding risk scenario displays the opposite pattern (see orange curve). Restoration investments are low initially due to the lack of an immediate threat. The manager opts to defer restoration costs to the future, but ramps up investment as sea-level rise increases flooding risks even from small storms. The total amount of restoration is similar under the two scenarios, but the timing is markedly different.

Figure 4 presents optimal planned restoration, meaning that decisions after period 0, $\{y_1, y_2, ..., y_T\}$, are

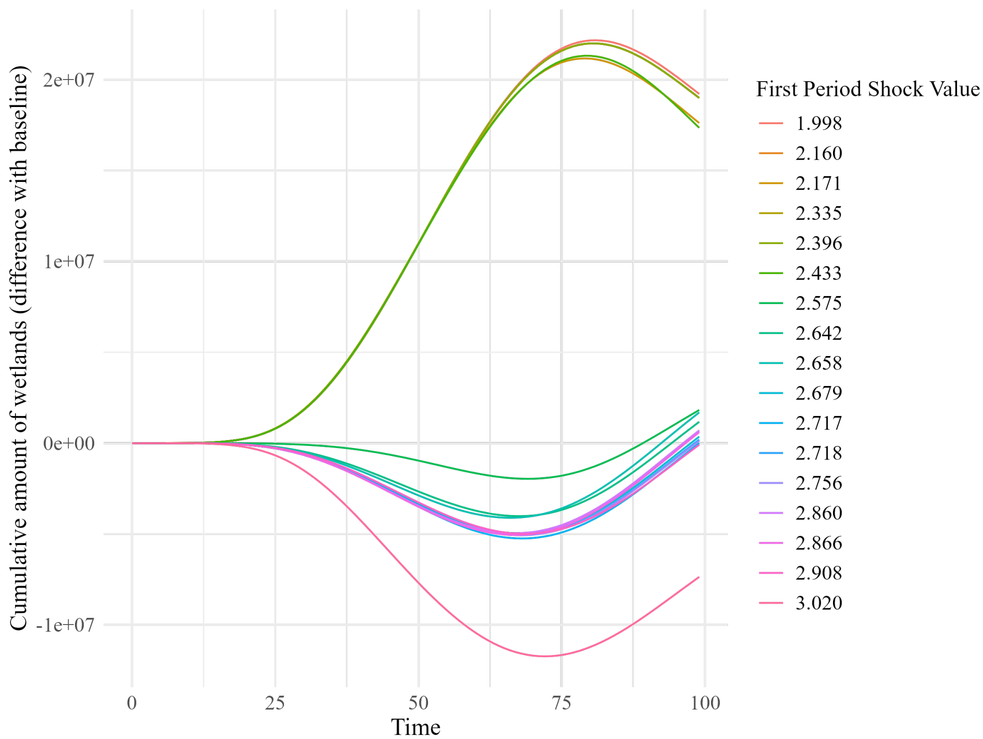

Figure 5: Departure from Baseline Plan Following Realization of First Period Shock

optimal in expectation given information available at time 0. It is instructive to consider how these values change as time progresses and new information becomes available. In Figure 5, we start the simulation in period 1 and vary the magnitude of the period 0 run-up value R_0. In contrast to Figure 4, the manager now chooses the value of y_1 with knowledge of $\mathbf{z}_0$, X_1, and $\mathbf{a}_1$. We plot the difference in the expected wetlands stock over time, X_t, relative to the baseline.[13] Differential effects of period 0 shocks on decisions to abandon and maintain properties translate into asymmetric and persistent differences in the optimally restored stock of wetlands. Larger flooding shocks at the end of period 0 result in more damaged properties that are optimally abandoned. Therefore, investments made early in the planning horizon are smaller and the stock of wetlands remains lower over the entire planning horizon. However, as sea level rises and the remaining properties are exposed to flooding risk, wetland stocks are built up. Thus, for a mid-range of period 0 shocks, there is little difference with the baseline by the end of the planning horizon. When period 0 shocks are small, properties sustain minimal damage in period 0 and it is optimal to maintain larger stocks of wetlands. The time path of the wetlands stock is similar for a range of small shocks because there is little difference among shocks in the period 0 property damage.

The analysis in Figure 5 can be extended to include realizations of the flooding shock over the entire

[13]Because the value of y_0 is chosen before R_0 is revealed, X_1 is the same for every realization of R_0 and equal to the baseline value.

planning horizon. We conduct a Monte Carlo simulation in which we draw R values for all periods (i.e., $\{R_0, R_1, ..., R_T\}$). For each of 25 sets of R values, we compute the optimal restoration choices in each period. In contrast to results presented above, these are the *actual* values of y_t^* based on realizations of $\mathbf{z}_t$, rather than the planned choices. In Appendix Figure 2, we plot for each set of R values the difference between the actual values of y_t^* and the planned values under the baseline scenario. There are large initial differences in investments as large (small) flooding shocks decrease (increase) the benefits of early restoration. However, by about year 60, we start to see convergence in the values of y_t^* as the future stock of wetlands is increasingly dependent on the growth of patches established earlier, and the marginal value of new investments is diminished. At this point, there is less incentive to respond to flooding shocks and the investment time paths "smooth out" as the end of the planning horizon approaches.

Sensitivity analysis is used to explore how optimal restoration varies with growth and cost parameters (Figure 6). For the growth function $f(\alpha) = e^{-\tau\alpha}$, we raise and lower the value of τ from its baseline value of 0.08. With faster growth ($\tau = 0.06$), initial investment falls dramatically relative to the baseline, but exceeds baseline investment in later years as the value of establishing new wetlands remains high even at the end of the planning horizon. With slower growth ($\tau = 0.12$), investment in restoration has much less value and is lower than baseline investment in all but the first few years. For the cost function $C(y_t) = \gamma y_t^2$, we raise and lower the value of γ from its baseline value of 1. When costs are high ($\gamma = 1.5$), investment is shifted from early in the planning horizon to later in order to diminish the contribution of costs to the present value of net benefits. The opposite pattern is seen with low costs ($\gamma = 0.5$).

A central motivation for the empirical application is to explore the effects of irreversibility and uncertainty on optimal restoration. We assume that once the manager decides not to repair a property following a flood, this abandonment decision cannot be reversed in the future. When combined with uncertainty over future flooding shocks, there may be value in delaying the irreversible abandonment decision in order to preserve the option to maintain or abandon a property in the future. For each property i, we compute the real option value (Dixit and Pindyck, 1994) as the difference between the value of $E_z[V(X_0, a_{i0}, z_{i0}]$ found using Bellman's equation in (15) and the expected net present value of the property in period 0 ($ENPV_0$) when the abandon/maintain decision is made using an $ENPV$ rule.[14] We find large option values associated with maintaining properties facing flooding risk. The mean (median) option value is \$221,832 (\$182,470), which is about 12% of the total value of properties in our sample. In aggregate, the option value is over \$1.6 billion. Appendix Figure 3 shows that restoration investments are higher under the Bellman rule than the $ENPV$

[14]Specifically, in every period t we compute the $ENPV$ of maintaining the property from t until the end of the planning horizon. As long as the $ENPV$ is positive, the property is maintained; otherwise it is abandoned. $ENPV_0$ is the discounted sum of rents minus expected damages from period 0 until the property is abandoned according to the $ENPV$ rule. The calculation of $E_z[V(X_0, a_{i0}, z_{i0}]$ and $ENPV_0$ are made at the optimal values of y_t; i.e., $\mathbf{y}_0^*(\mathbf{a}_0, X_0)$, the solution to (16).

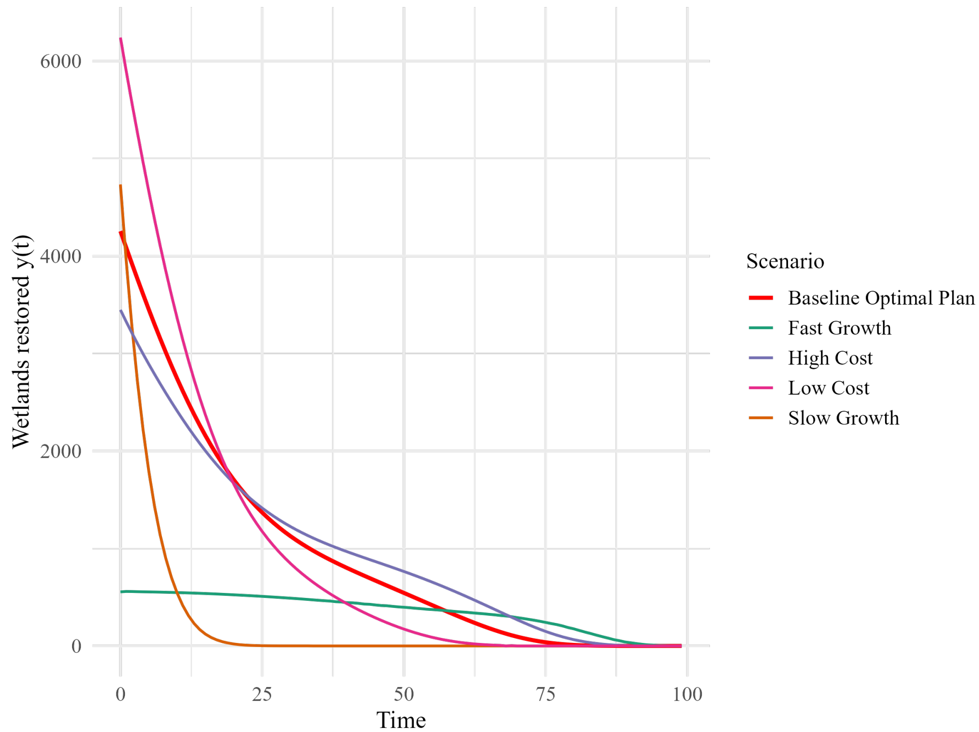

Figure 6: Sensitivity of Results to Cost and Growth Assumptions

rule. The $ENPV$ rule abandons properties too early, which mutes the incentive to invest in restoration. When option values are accounted for in abandonment decisions, properties are more valuable and there is greater incentive to protect them from flooding.

Table 1: Option Value of Maintaining Properties Subject to Flooding Risks

	Value
Mean Option Value	$221,832
Median Option Value	$182,470
Percent of Total Property Value	11.94%
Total Value of Options	$1,628,500,000

5 Discussion and Conclusions

Although preservation of unique environments is critical for the provision of ecosystem services, there is growing recognition of the need to restore damaged and degraded ecosystems. Previous research has addressed some of the economic dimensions of ecosystem restoration, but missing from the literature is a full treatment of the dynamic tradeoffs. Dynamics are fundamental to the problem because of growth in restored areas and environmental change. In this paper, we study the optimal timing and extent of investments in ecosystem restoration when damages from un-restored systems are non-stationary. Our analysis applies to a wide range of restoration activities, including planting of native species in forest, grassland, and marine environments, as well as changes in land management practices.

We begin by developing an optimal control model in which the manager chooses the amount of restoration in each period to maximize the present value of net benefits from the ecosystem. The key features of the model, which depart from traditional specifications in resource economics, are a state variable that is time- and age-dependent and non-stationary damages. An age-dependent growth model is needed to accurately characterize a large range of restoration activities. Although the resulting state equation is a partial differential equation, we obtain an explicit solution to the model by using a boundary Hamiltonian and the Method of Characteristics. An additional contribution of this study is to show how a broader class of optimal resource models can be solved by leveraging insights from vintage capital models.

Our key analytical result is an expression for the shadow value of the state variable $x(t, \alpha)$. The value of a marginal increase in an age α restored patch in time t is the present value of the avoided damage over the lifetime of the patch, accounting for its growth over time and evolving marginal damages. In each period t, restoration should continue up until the marginal costs of restoration are equal to the shadow value. Numerical simulations show that optimal investment occurs early in the planning horizon. Even though damages are low initially, early investments are made to take advantage of growth in restored patches so that future damages are mitigated. Only when damages are far in the future is it optimal to ramp up investment over time. When growth rates are held constant, we remove the incentive to establish new patches with higher growth. In contrast to a model with age-dependent growth, investment declines rapidly rather than being maintained at a high level over time.

The formulation of the optimal control model assumes that the manager has selected a particular restoration project. The solution to the control problem answers the question of how to implement it optimally. A separate question is whether to pursue that particular project, an alternative project, or, indeed, any project at all. Although we do not answer this question directly, we note that the control problem in (2) is a necessary ingredient to address it. Specifically, the objective function in (2), evaluated at the optimal solution, gives the maximized present value of net benefits from the project under consideration. If the net benefits are negative, then it is inadmissible according to cost-benefit criteria. If positive, then the project's net benefits can be compared to those generated by alternative projects.

To further explore the dynamics of optimal restoration, we apply the model to coastal wetlands restoration. Coastal wetlands protect real estate from the combined effects of sea-level rise and storm surge, while providing a range of ecosystem services and avoiding the environmental drawbacks of coastal hardening. In our model, future storm surge is stochastic and damages are non-stationary due to sea-level rise. We allow for irreversible abandonment of properties, which when combined with uncertainty over damages gives rise to potential option value. We expect option value to be important in settings where large shocks to an ecosystem make it optimal to abandon restoration activities, at least in severely-affected areas. Examples include marine heat waves that cause mass coral bleaching and establishment of invasive species that are extremely difficult to remove. In these cases, it may be optimal to continue restoration in order to preserve the option to abandon or continue in the future. Although we treat abandonment of damaged coastal properties as irreversible, future research could explore the conditions under which rebuilding in abandoned locations is warranted.[15]

Consistent with the analytical model, we find that planned investment in coastal wetlands occurs early in the planning period to take advantage of growth in restored patches. Sea-level rise magnifies this effect by raising expected future damages. The time profile of planned investment is sensitive to the severity of damages. When they are expected to be high, more investment occurs early to build up the stock of wetlands. However, investment in new patches can eventually decrease as growth in the established stock provides protection against severe storms. The opposite pattern is seen when damages are expected to be low in the future. In this case, investments can be delayed, lowering the present value of restoration costs. The size of the future stock is lower since growth is diminished, necessitating a higher level of future investment than when storms are expected to be severe. In our study, we assumed a stationary storm surge distribution based on historical data. If climate change increases the severity of future storms, one could allow for a shifting distribution over time (Tebaldi et al., 2012), which would likely increase incentives for early restoration investments.

As time unfolds and damages occur, adjustments are made in planned restoration investments. In particular, when flooding damages are large early in the planning horizon, some properties are optimally abandoned, which decreases the benefits of further investments. As a consequence, the stock of restored wetlands is lower throughout the planning horizon. In our model, the current shock does not provide information about future shocks. As such, high flooding damages today do not affect expectations about damages in the future. An interesting extension of the analysis would be to allow expectations about future storm surges to be updated based on the current realizations.

A key motivation for including uncertainty in the empirical application is to gauge the importance of option value. In our setting, option value is the value to the manager of having the flexibility to delay irreversible abandonment decisions. We find that option values are large. We estimate that the average option value in Huntington Beach, CA, is approximately \$200,000 per property, and more than \$1.6 billion in the aggregate. The broader implication of this finding is that the net benefits from restoration may be

[15] Chapter 7 in Dixit and Pindyck (1994) offers a starting point for this analysis.

enhanced significantly if flexible decision making is allowed. Often, there are substantial time lags between the proposal and implementation stages of restoration projects. For example, under the Collaborative Forest Landscape Restoration Program[16] administered by the U.S. Forest Service, communities propose specific forest restoration projects. This involves a lengthy application process in addition to potential delays due to National Environmental Policy Act (NEPA) requirements. If the location, extent, and timing of restoration can adapt to changing environmental and economic conditions, society will obtain greater value from the resulting ecosystem service flows.

References

Arkema, K. K., Guannel, G., Verutes, G., Wood, S. A., Guerry, A., Ruckelshaus, M., Kareiva, P., Lacayo, M., and Silver, J. M. (2013). Coastal habitats shield people and property from sea-level rise and storms. *Nature climate change*, 3(10):913–918.

Arneth, A., Olsson, L., Cowie, A., Erb, K.-H., Hurlbert, M., Kurz, W. A., Mirzabaev, A., and Rounsevell, M. D. (2021). Restoring degraded lands. *Annual Review of Environment and Resources*, 46(1):569–599.

Benhabib, J. and Rustichini, A. (1991). Vintage capital, investment, and growth. *Journal of economic theory*, 55(2):323–339.

Blignaut, J., Aronson, J., and de Wit, M. (2014). The economics of restoration: looking back and leaping forward. *Annals of the New York Academy of Sciences*, 1322(1):35–47.

Caputo, M. R. and Wilen, J. E. (1995). Optimal cleanup of hazardous wastes. *International Economic Review*, pages 217–243.

Cattaneo, M. D., Crump, R. K., Farrell, M. H., and Feng, Y. (2024). Nonlinear binscatter methods. *arXiv preprint arXiv:2407.15276*.

Davis, M. A., Larson, W. D., Oliner, S. D., and Shui, J. (2021). The price of residential land for counties, zip codes, and census tracts in the united states. *Journal of Monetary Economics*, 118:413–431.

Davis, S. and Skaggs, L. (1992). Catalog of residential depth-damage functions. *US Army Corps of Engineers, Institute for Water Resources, Ft. Belvoir, VA.*

Dixit, A. K. and Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton university press.

[16] https://www.fs.usda.gov/restoration/CFLRP/

Erikson, L. H., Espejo, A., Barnard, P. L., Serafin, K. A., Hegermiller, C. A., O'Neill, A., Ruggiero, P., Limber, P. W., and Mendez, F. J. (2018). Identification of storm events and contiguous coastal sections for deterministic modeling of extreme coastal flood events in response to climate change. *Coastal Engineering*, 140:316–330.

Fairchild, T. P., Bennett, W. G., Smith, G., Day, B., Skov, M. W., Möller, I., Beaumont, N., Karunarathna, H., and Griffin, J. N. (2021). Coastal wetlands mitigate storm flooding and associated costs in estuaries. *Environmental Research Letters*, 16(7):074034.

Feichtinger, G., Hartl, R. F., Kort, P. M., and Veliov, V. M. (2006). Capital accumulation under technological progress and learning: a vintage capital approach. *European Journal of Operational Research*, 172(1):293–310.

Goetz, R.-U., Hritonenko, N., and Yatsenko, Y. (2008). The optimal economic lifetime of vintage capital in the presence of operating costs, technological progress, and learning. *Journal of Economic Dynamics and Control*, 32(9):3032–3053.

Grossinger, R., Stein, E. D., Cayce, K., Askevold, R., Dark, S., and Whipple, A. (2011). Historical wetlands of the southern california coast. *SFEI Contribution*, 596.

Hartman, R. (2018). The harvesting decision when a standing forest has value. In *Economics of Forestry*, pages 295–302. Routledge.

Kulp, S. A. and Strauss, B. H. (2019). New elevation data triple estimates of global vulnerability to sea-level rise and coastal flooding. *Nature communications*, 10(1):1–12.

Lappi, P. (2018). Optimal clean-up of polluted sites. *Resource and Energy Economics*, 54:53–68.

Mirzabaev, A. and Wuepper, D. (2023). Economics of ecosystem restoration. *Annual Review of Resource Economics*, 15(1):329–350.

Narayan, S., Beck, M. W., Reguero, B. G., Losada, I. J., Van Wesenbeeck, B., Pontee, N., Sanchirico, J. N., Ingram, J. C., Lange, G.-M., and Burks-Copes, K. A. (2016). The effectiveness, costs and coastal protection benefits of natural and nature-based defences. *PloS one*, 11(5):e0154735.

Neeson, T. M., Ferris, M. C., Diebel, M. W., Doran, P. J., O'Hanley, J. R., and McIntyre, P. B. (2015). Enhancing ecosystem restoration efficiency through spatial and temporal coordination. *Proceedings of the National Academy of Sciences*, 112(19):6236–6241.

Paine, C. T., Marthews, T. R., Vogt, D. R., Purves, D., Rees, M., Hector, A., and Turnbull, L. A. (2012). How to fit nonlinear plant growth models and calculate growth rates: an update for ecologists. *Methods in Ecology and Evolution*, 3(2):245–256.

Reimann, L., Vafeidis, A. T., and Honsel, L. E. (2023). Population development as a driver of coastal risk: Current trends and future pathways. *Cambridge Prisms: Coastal Futures*, 1:e14.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033.

Smith, M. D., Slott, J. M., McNamara, D., and Murray, A. B. (2009). Beach nourishment as a dynamic capital accumulation problem. *Journal of Environmental Economics and Management*, 58(1):58–71.

Stein, E. D., Cayce, K., Salomon, M., Bram, D. L., De Mello, D., Grossinger, R., and Dark, S. (2014). Wetlands of the southern california coast: Historical extent and change over time. *Southern California coastal water research project technical report*, 826:1–50.

Sun, F. and Carson, R. T. (2020). Coastal wetlands reduce property damage during tropical cyclones. *Proceedings of the National Academy of Sciences*, 117(11):5719–5725.

Swanson, N. R. and White, H. (1997). A model selection approach to real-time macroeconomic forecasting using linear models and artificial neural networks. *Review of Economics and Statistics*, 79(4):540–550.

Syphard, A. D., Brennan, T. J., Rustigian-Romsos, H., and Keeley, J. E. (2022). Fire-driven vegetation type conversion in southern california. *Ecological Applications*, 32(6):e2626.

Tebaldi, C., Strauss, B. H., and Zervas, C. E. (2012). Modelling sea level rise impacts on storm surges along us coasts. *Environmental Research Letters*, 7(1):014032.

Wilson, E. O. (1999). *The diversity of life*. WW Norton & Company.

Xepapadeas, A. and de Zeeuw, A. (1999). Environmental policy and competitiveness: the porter hypothesis and the composition of capital. *Journal of environmental economics and Management*, 37(2):165–182.

A Appendices

A.1 The State Equation and Aggregate Restoration

In time t, the aggregate amount of restoration is given by $X(t) = \int_0^t x(t,\alpha)d\alpha$. We show here that the state equation can be rewritten as a function of aggregate restoration when $f(\alpha)$ is a constant; however, we are unable to simplify the state equation in this way when $f(\alpha)$ is a concave function. Taking the time derivative of $X(t)$, substituting the state equation in (2), and solving the second integral yields:

$$\begin{aligned}
\dot{X}(t) &= x(t,t) + \int_0^t \frac{\partial x(t,\alpha)}{\partial t} d\alpha \\
&= x(t,t) + \int_0^t [f(\alpha)x(t,\alpha) - \frac{\partial x(t,\alpha)}{\partial \alpha}] d\alpha \\
&= x(t,t) + \int_0^t f(\alpha)x(t,\alpha)d\alpha - [x(t,t) - x(t,0)] \\
&= x(t,0) + \int_0^t f(\alpha)x(t,\alpha)d\alpha
\end{aligned} \tag{20}$$

$x(t,0)$ is equal to new additions to the aggregate stock (e.g., from investments in restoration). Suppose $f(\alpha)$ takes the form $f(\alpha) = \theta$, where θ is a strictly positive parameter. In this case, (20) becomes:

$$\begin{aligned}
\dot{X}(t) &= x(t,0) + \int_0^t \theta x(t,\alpha)d\alpha \\
&= x(t,0) + \theta X(t)
\end{aligned} \tag{21}$$

When $f(\alpha)$ depends on α, we are unable to simplify the integral in (20) and write $\dot{X}(t)$ in terms of $X(t)$. For example, applying integration by parts to (20) does not yield a function of $X(t)$.

When there are no new additions to the stock, $x(t,0) = 0$, and equation (21) implies that the growth rate in the aggregate stock is constant:

$$\frac{\dot{X}}{X} = \theta \tag{22}$$

A standard specification for $\dot{X}$ in the renewable resource economics literature is:

$$\begin{aligned}
\dot{X} &= F(X) \\
&= rX\left(1 - \frac{X}{K}\right)
\end{aligned} \tag{23}$$

where r is the intrinsic growth rate and K is the carrying capacity. In this model, the growth rate in the stock is given by:

$$\frac{\dot{X}}{X} = r - r\frac{X}{K} \tag{24}$$

The growth rate reaches a maximum of r at $X = 0$ and declines to 0 at $X = K$.

A.2 Analytical Solution

A.2.1 Manager's Problem

$$\begin{gathered} \underset{y(t)}{max} \int_0^T e^{-rt}\big[\int_0^t b(t) - p(t)(\bar{x} - x(t,\alpha))d\alpha - C(y(t))\big]dt \quad \text{s.t.} \\ \frac{\partial x(t,\alpha)}{\partial t} + \frac{\partial x(t,\alpha)}{\partial \alpha} = x(t,\alpha)f(\alpha) \quad \alpha \in [0,t] \quad x(t,0) = y(t) \quad \lambda(T,\alpha) = 0 \end{gathered} \tag{25}$$

To be clear about definitions, α is the age or years since a patch was restored, t is the years since the start of our timeline, and $t - \alpha$ is the birthday or year restored for a specific patch.

A.2.2 Hamiltonian and Solution

Current Value Hamiltonian

$$\tilde{H}(x,y,\lambda) = b(t) - p(t)(\bar{x} - x(t)) + \lambda(\alpha,t)f(\alpha)x(\alpha,t) \tag{26}$$

Boundary Hamiltonian

$$H_0 = -C(y(t)) + \lambda(0,t)y(t) \tag{27}$$

First-Order Conditions

$$\frac{\partial \tilde{H}}{\partial y} = 0 \tag{28}$$

$$\frac{\partial H_0}{\partial y} = 0 = -C'(y(t)) + \lambda(t,0) \tag{29}$$

Adjoint Equation

$$\frac{\partial \lambda}{\partial t} + \frac{\partial \lambda}{\partial \alpha} = r\lambda - \frac{\partial \tilde{H}}{\partial x} = -p(t) + \lambda(r - f(\alpha)) \tag{30}$$

A.2.3 Method of Characteristics

In this section, we present the method of characteristics and show how it applies to our problem. The first step is to assume t and α are functions of s, and find how the shadow price λ changes with s. From the chain rule, we have:

$$\frac{\partial \lambda(\alpha,t)}{\partial s} = \frac{\partial \lambda(\alpha,t)}{\partial t} * \frac{\partial t}{\partial s} + \frac{\partial \lambda(\alpha,t)}{\partial \alpha} * \frac{\partial \alpha}{\partial s}$$

This expression matches the left-hand side of the adjoint equation in (30) if $\frac{\partial t}{\partial s} = \frac{\partial \alpha}{\partial s} = 1$. Thus,

$$\frac{\partial \lambda}{\partial s} = -p(t) + \lambda(r - f(\alpha)) \tag{31}$$

Then, if we solve for ds in the three expressions and set them equal, we have:

$$\partial t = \partial \alpha = \frac{\partial \lambda}{-p(t) + \lambda(r - f(\alpha))}$$

Rearranging yields:

$$\begin{aligned}\frac{\partial \lambda(t,\alpha)}{\partial t} &= -p(t) + \lambda(t,\alpha)(r - f(\alpha)) \\ \frac{\partial \lambda(t,\alpha)}{\partial t} - \lambda(t,\alpha)(r - f(\alpha)) &= -p(t)\end{aligned}$$

The second equation is an ODE that we can now solve. We start by multiplying through by $\mu(t,\alpha)$:

$$\frac{\partial \lambda(t,\alpha)}{\partial t}\mu(t,\alpha) - \lambda(t,\alpha)(r - f(\alpha))\mu(t,\alpha) = -p(t)\mu(t,\alpha)$$

Then, setting the integrating equation to $\frac{\partial \mu(t,\alpha)}{\partial t} = -(r - f(\alpha))\mu(t,\alpha)$ and substituting gives:

$$\frac{\partial \lambda(t,\alpha)}{\partial t}\mu(t,\alpha) + \frac{\partial \mu(t,\alpha)}{\partial t}\lambda(t,\alpha) = -\mu(t,\alpha)p(t)$$

Integrating both sides from t to T yields:

$$\lambda(s,\alpha)\mu(s,\alpha)|_t^T = -\int_t^T \mu(s,\alpha)p(s)ds \tag{32}$$

Before we can solve (32), we have to solve for the integrating equation:

$$\begin{aligned}\frac{\partial \mu(t,\alpha)}{\partial t} &= -(r - f(\alpha))\mu(t) \\ \frac{\frac{\partial \mu(t,\alpha)}{\partial t}}{\mu(t,\alpha)} &= -(r - f(\alpha)) \\ ln(\mu(s,\alpha)) &= -\int_s^T (r - f(\alpha + \rho - t))d\rho \\ \mu(s,\alpha) &= e^{-\int_s^T (r - f(\alpha+\rho-t))d\rho}\end{aligned}$$

Now, we solve the left-hand side of (32) to obtain:

$$\begin{aligned}\lambda(s,\alpha)\mu(s,\alpha)|_t^T = \lambda(T,\alpha)\mu(T,\alpha) - \lambda(t,\alpha)\mu(t,\alpha) = 0 - \lambda(t,\alpha)e^{-\int_t^T (r-f(\alpha+\rho-t))d\rho} \\ = -\lambda(t,\alpha)e^{-\int_t^T (r-f(\alpha+\rho-t))d\rho}\end{aligned}$$

This result makes use of the boundary conditions $\lambda(T,\alpha) = 0$ and the integrating equation evaluated at t. Substituting the integrating equation into the right-hand side of (32), we arrive finally at equation (8):

$$\lambda(t,\alpha) = e^{\int_t^T (r-f(\alpha+\rho-t))d\rho} \int_t^T e^{-\int_s^T (r-f(\alpha+\rho-t))d\rho} p(s)ds = \int_t^T e^{-\int_t^s (r-f(\alpha+\rho-t))d\rho} p(s)ds$$

Since the discount rate is constant, we can alternatively write the shadow price as:

$$\lambda(t,\alpha) = \int_t^T e^{-r(s-t)} e^{\int_t^s f(\alpha+\rho-t))d\rho} p(s)ds$$

Using the first-order condition for the boundary Hamiltonian, we can recover $y(t)$ from:

$$C'(y(t)) = \int_t^T e^{-r(s-t)} e^{\int_t^s f(\rho-t))d\rho} p(s)ds$$

Finally, we derive an expression for $x(t,\alpha)$ using the state equation:

$$\frac{\partial x(t,\alpha)}{\partial t} + \frac{\partial x(t,\alpha)}{\partial \alpha} = x(t,\alpha)f(\alpha)$$

We can once again apply the Method of Characteristics, which gives us:

$$\begin{aligned}\partial t = \partial \alpha = \frac{\partial x(t,\alpha)}{x(t,\alpha)f(\alpha)} \\ \frac{\partial x(t,\alpha)}{\partial \alpha} = x(t,\alpha)f(\alpha) \\ ln(x(t,\alpha)) - ln(x(t-\alpha,0)) = \int_0^\alpha f(\alpha)d\alpha \\ x(t,\alpha) = x(t-\alpha,0)e^{\int_0^\alpha f(\alpha)d\alpha}\end{aligned}$$

This is equation (10). Alternatively, we can apply the initial condition $x(t,0) = y(t)$ to obtain:

$$x(t,\alpha) = y(t-\alpha)e^{\int_0^\alpha f(\alpha)\partial\alpha}$$

This result shows that the total amount of restoration of age α in time t is the amount restored α years ago, adjusted for growth over the subsequent α years.

A.3 Functional Forms to Illustrate the Analytical Solution

We adopt the following functional forms to produce Figures 2 and 3:

$$C(y(t)) = (y(t) - 1)^2$$
$$f(\alpha) = \frac{1}{\alpha + 1}$$
$$p(t) = \frac{N}{1 + e^{-B(t-A)}}$$

Parameter values are:

Base: $N = 10 \quad r = 0.03 \quad A = 5 \quad B = 0.75 \quad T = 100$

Far Damage: $N = 10 \quad r = 0.03 \quad A = 40 \quad B = 0.75 \quad T = 100$

High Damage: $N = 15 \quad r = 0.03 \quad A = 5 \quad B = 0.75 \quad T = 100$

High Interest Rate: $N = 10 \quad r = 0.05 \quad A = 5 \quad B = 0.75 \quad T = 100$

A.4 Equivalence of Solution Method to Backwards Induction

We show that the solution to the problem in (16) is equivalent to what is obtained with backwards induction. We demonstrate this using a two-period model for one property. It is straightforward, though notationally cumbersome, to extend the results to longer time horizons and multiple properties.

To help provide intuition for the solution, Figure (1) shows a portion of the decision tree for the problem. For ease of illustration, we assume in the figure, but not in the results derived below, that in each time period $t = \{0, 1\}$ restoration is a binary decision $y_t = \{y_{t1}, y_{t2}\}$, the shock takes two values $z_t = \{z_{t1}, z_{t2}\}$, and, as in the original problem, maintain/abandon is a binary decision $d_t = \{d_{t1}, d_{t2}\}$. Figure (1) shows the $t = 1$ decision tree conditional on choices y_{01} and d_{01} in period 0 and the shock z_{01}. The decision-maker first chooses y_1, the shock z_1 happens, and then the maintain/abandon decision d_1 is made. This results in payoffs m_1 that are dependent on the full set of time 0 and 1 choices and shocks. To focus on the choices by the decision-maker, we write the payoffs as a function of the full sequence of controls rather than the state variables.[17]

[17] We could equivalently write the payoffs in Figure (1) as a function of state variables X and a rather than the period 0 choices and shock. Note that the payoff function $m(\cdot)$ captures both the net rents in (14) and the costs of restoration in (16).

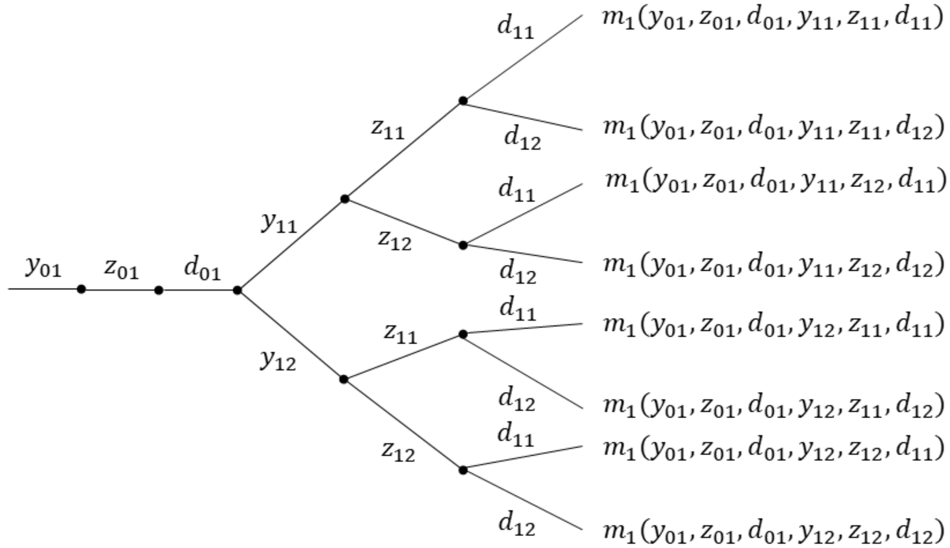

Figure 1: A Portion of the Decision Tree for the Restoration and Maintain/Abandon Problem

To find the backwards induction solution, we start at the right of the decision tree and, from each node (indicated by a dot), choose the value of d_1 that results in the highest payoff. In general, we solve:

$$\max_{d_1} m_1(y_0, z_0, d_0, y_1, z_1, d_1) \tag{33}$$

which yields the maximized payoff at any node as:

$$m_1(y_0, z_0, d_0, y_1, z_1, d_1^*) \tag{34}$$

where d_1^* maximizes m_1 for given values of y_0, z_0, d_0, y_1 and z_1. Moving left in the decision tree, z_1 is a random variable chosen by "nature" rather than the decision-maker. Thus, at the node preceding the shock we have the expected payoff:

$$E_{z_1} m_1(y_0, z_0, d_0, y_1, z_1, d_1^*) \tag{35}$$

Because z_0 and z_1 are independent, the expectation is not conditioned on z_0. Finally, y_1 is chosen to maximize the expected payoff:

$$\max_{y_1} E_{z_1} m_1(y_0, z_0, d_0, y_1, z_1, d_1^*) \tag{36}$$

which yields the maximized expected payoff for period 1:

$$E_{z_1} m_1(y_0, z_0, d_0, y_1^*, z_1, d_1^*) \tag{37}$$

where y_1^* is the optimal value of y_1.

The period 0 decisions and shock yield the payoff $m_0(y_0, z_0, d_0)$, which is added to the period 1 expected payoff to obtain:

$$m_0(y_0, z_0, d_0) + E_{z_1} m_1(y_0, z_0, d_0, y_1^*, z_1, d_1^*) \tag{38}$$

where the discount factor is set to 1. Following the same steps as for period 1, we obtain at the starting node for period 0 the maximized expected value (EV^*) for the problem:

$$EV^* = E_{z_0}[m_0(y_0^*, z_0, d_0^*) + E_{z_1} m_1(y_0^*, z_0, d_0^*, y_1^*, z_1, d_1^*)] \tag{39}$$

$$= E_{z_0} m_0(y_0^*, z_0, d_0^*) + E_z m_1(y_0^*, z_0, d_0^*, y_1^*, z_1, d_1^*) \tag{40}$$

where E_z denotes the joint expectation for z_0 and z_1.

We show, next, that we obtain the same EV^* using the solution method in (16). This method begins by fixing y_0 and y_1 and then applying backwards induction to obtain d_1^*. The initial steps are the same as above, except that we do not apply the maximization in (36). Thus, the period 0 payoff in (38) is:

$$m_0(y_0, z_0, d_0) + E_{z_1} m_1(y_0, z_0, d_0, y_1, z_1, d_1^*) \tag{41}$$

where y_1 is the pre-set, rather than the optimal value. When we arrive at the starting node for period 0 we have:

$$EV = E_{z_0} m_0(y_0, z_0, d_0^*) + E_z m_1(y_0, z_0, d_0^*, y_1, z_1, d_1^*) \tag{42}$$

Equation (42) gives the maximized expected value conditional on y_0 and y_1. The next step is to find the optimal values of y_0 and y_1 by solving:

$$\max_{y_0, y_1} EV = E_{z_0} m_0(y_0, z_0, d_0^*) + E_z m_1(y_0, z_0, d_0^*, y_1, z_1, d_1^*) \tag{43}$$

The maximized objective function is:

$$EV^* = E_{z_0} m_0(y_0^*, z_0, d_0^*) + E_z m_1(y_0^*, z_0, d_0^*, y_1^*, z_1, d_1^*) \tag{44}$$

Equations (44) and (39) are equivalent and, thus, we have established that the method in (16) yields the same solution as backwards induction.

A.5 Hedonic Rent Equation

Annual rental values for unavailable for approximately 19% of the properties in our sample. Rather than omit these properties, we predict rents using a standard hedonic price model fitted to the observations for which we have complete data. The dependent variable is $log(rent)$, the natural log of estimated rent values web-scraped from Redfin. For independent variables, we include structural characteristics such as the number of bedrooms, number of bathrooms, square footage, and year built. These variables are taken from Redfin and have minimal missing observations. We also include the distance to the coast to capture amenity values associated with beach access. Finally, we control for area demographics, drawn from the 2020 Census at the Census tract level. The Redfin data in our sample covers four Census tracts.

Estimation results are given in the table below. Although the purpose of the regression model is prediction and not estimation of individual coefficients, we note that most of the estimates are significantly different from zero and have the expected signs. An exception is the coefficient on distance to coast, which is imprecisely estimated due to a lack of variation within the sample.

Estimates from Hedonic Regression Model

Variable	**Estimate**	**Std. Error**	**p-value**
Intercept	13.1500	0.4225	$<$ 2e-16 ***
Distance to Coast (m)	5.266e-07	1.797e-06	0.7695
Square Footage	0.0002367	7.675e-06	$<$ 2e-16 ***
Beds	0.07840	0.004666	$<$ 2e-16 ***
Baths	0.09403	0.007211	$<$ 2e-16 ***
Year Built	-0.002757	0.0002108	$<$ 2e-16 ***
Median Age	-0.01050	0.0004607	$<$ 2e-16 ***
Percent Black	-0.0003091	6.018e-05	2.88e-07 ***
Percent Asian	-3.952e-05	1.028e-05	0.000122 ***
Median Income	1.689e-06	9.953e-08	$<$ 2e-16 ***
Average Education	7.836e-05	1.081e-05	4.64e-13 ***
Unemployment Rate	-0.0003939	5.599e-05	2.17e-12 ***

Note: Residual standard error = 0.2712 on 7107 degrees of freedom; Multiple R^2 = 0.7097; Adjusted R^2 = 0.7093; F-statistic = 1580 on 11 and 7107 DF; p-value $<$ 2.2e-16. Significance codes: *** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$.

A.6 Additional Numerical Simulation Figures

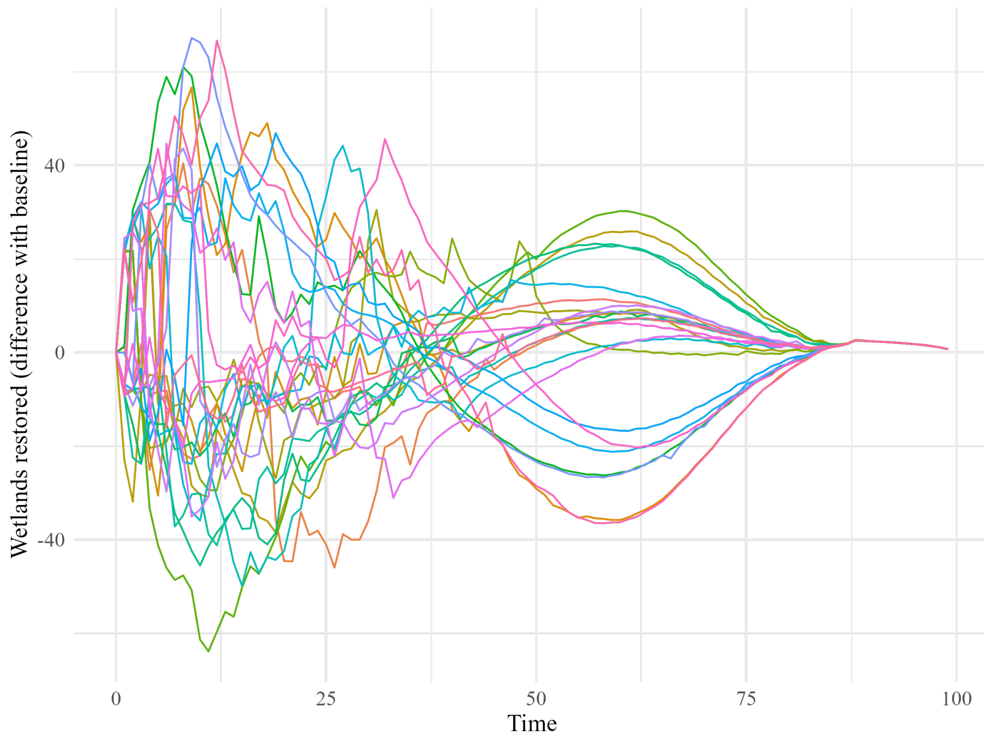

Figure 2: Difference in Wetlands Restored Between Monte Carlo Simulated Values of R_t and Baseline Scenario

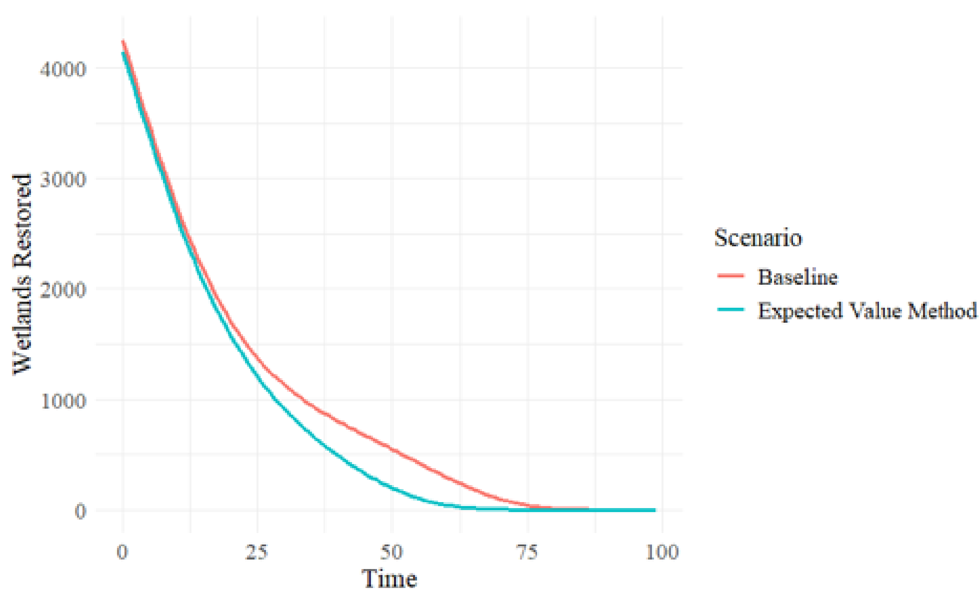

Figure 3: Wetlands Restored When Abandon and Maintain Decisions are Made Using Bellman's Equation in (15) (Baseline) and an ENPV Rule (see footnote 14)