

Balancing multiple interventions for dynamically efficient kelp forest restoration under marine heatwave uncertainty

Kyumin Kim¹, Marissa Baskett², Michael Springborn²

¹Department of Agricultural and Resource Economics, University of California, Davis, CA

²Department of Environmental Science and Policy, University of California, Davis, CA

Abstract

Kelp forests are in decline globally, threatening critical socio-ecological functions such as marine species refugia, coastal productivity, and community livelihoods. In Northern California, the collapse of kelp forests in the last decade, driven by marine heatwaves (MHWs) and a surge in kelp-consuming purple sea urchins, has caused severe ecological degradation and economic hardship. In this paper we develop the first bioeconomic model to evaluate dynamic kelp forest restoration strategies. The model integrates (i) kelp-urchin population feedback dynamics, (ii) two restoration options—kelp outplanting and purple urchin culling—that can be varied in intensity, (iii) restoration costs and economic benefits, including both fishery and direct kelp values, and (iv) a stochastic MHW regime with uncertain frequency. Results show a sharp difference between the two restoration options in how they optimally respond to changes in the population levels, with outplanting used selectively and culling used consistently. Similarly, we find that the lost value from failing to make use of both restoration tools is asymmetric: use of outplanting alone results in substantial loss of value (around half) while losses from using urchin removal alone are much more modest (unless the kelp population is very low). These results are sensitive to the main environmental stress: if MHW frequency shifts permanently to the high level recently experienced, the optimal use of kelp outplanting broadens over possible conditions. Overall we find that the two restoration approaches play distinct, state-dependent roles in efficient restoration that respond dynamically to varying ecological and economic conditions.

Keywords: Bioeconomic modeling, Kelp forest restoration, Stochastic dynamic optimization, Decision-making under uncertainty, Marine heatwaves, Coupled human–natural systems

1 Introduction

Kelp forests are vital to coastal marine ecosystems, offering habitat, food, and nursery grounds for diverse marine life, including commercially important species such as lobster, abalone, and red sea urchin. They also provide non-market ecosystem services—such as water purification, coastal protection, nutrient regulation, and carbon sequestration—that support both biodiversity and human well-being (Hynes et al., 2021; Smale et al., 2016; Eger et al., 2020). However, kelp forests have experienced dramatic declines globally due to multiple stressors, including ocean warming, pollution such as eutrophication, and overfishing (Hynes et al., 2021; Fredriksen et al., 2020; Wernberg et al., 2018; Carnell et al., 2025).

In California, a record-breaking marine heatwave (MHW) from 2013 to 2015 produced warm, nutrient-poor waters that hindered kelp growth. Concurrently, an outbreak of sea star wasting disease in 2013 led to the functional extinction of sunflower sea stars (*Pycnopodia helianthoides*), key predators of kelp-grazing purple urchins (*Strongylocentrotus purpuratus*) (Harvell et al., 2019). The collapse of predation pressure, combined with ocean warming, sharply reduced drift kelp—detached fronds that serve as a primary food source for urchins, triggering a surge in purple urchin populations and intensifying grazing pressure on living kelp, including the overconsumption of adult stipes and even the holdfast (equivalent to stem and root of land plants). As a result, bull kelp (*Nereocystis luetkeana*) forest cover in Northern California declined by over 90% relative to pre-disturbance levels (Lorenzo and Mantua, 2016; Rogers-Bennett and Catton, 2019; McPherson et al., 2021).

The dramatic decline of kelp forests has also had significant socio-economic impacts on the welfare of coastal communities, operating through multiple channels. In 2017, kelp loss led to an estimated 80% mortality in the red abalone population, resulting in the fishery’s closure in 2018 (Rogers-Bennett and Catton, 2019). More than 35,000 recreational fishers had contributed an estimated \$44 million annually to local economies through tourism and other activities related to the red abalone fishery (Reid et al., 2016). Associated travel expenditures provided substantial economic value to California’s coastal communities (Rogers-Bennett and Catton, 2019). Kelp forests have yet to recover, and the timeline for reopening the red abalone fishery remains uncertain.

Similarly, kelp forests have supported stakeholders in the red sea urchin (*Mesocentrotus franciscanus*) fishery through commercial fishing operations. This includes urchin divers, local community members employed by seafood processing companies for urchin processing (extracting uni), and various other related jobs. Like the red abalone recreational fishery, the red sea urchin fishery previously generated \$2-3 million in annual fishing revenues for local communities. However, this fishery suffered a catastrophic decline due to kelp forest loss, resulting in an 80% revenue reduction and a federal disaster declaration in 2016-2017

(California Department of Fish and Wildlife, 2020; Ward et al., 2022). Fishers’ knowledge is formulated through experiential understanding of marine and freshwater environments that fish harvesters accumulate while operating in their respective fisheries (Hind, 2014). When fisheries become unavailable or are disrupted, this knowledge accumulation process that depends on active fishing practice is interrupted. Therefore, this decline may also cause the loss of important but intangible assets in coastal communities, such as specialized professional knowledge for this dive fishery. Similarly, indigenous communities in California have also been heavily impacted, as historically they relied on the kelp ecosystem for subsistence, medicine, and cultural ceremonies (California Ocean Protection Council, 2021).

Despite growing interest in kelp forest restoration, there remains limited normative research assessing how restoration should be implemented. In the economics literature, studies have focused on valuation and cost–benefit aspects. For instance, one strand of research has emphasized the valuation of kelp forest ecosystem services—such as nutrient cycling, carbon sequestration, and enhanced fisheries productivity—as well as existence values derived from willingness-to-pay estimates using surveys, choice experiments, or meta-analyses (Vásquez et al., 2014; Bennett et al., 2016; Hynes et al., 2021; Williams et al., 2022). These studies demonstrate substantial economic value associated with kelp ecosystems.

Beyond evaluating the economic value of kelp forests, two studies have assessed the restoration of these systems. Carnell et al. (2025) presented the first spatially explicit benefit–cost analysis of kelp forest restoration in southern Australia. The study evaluated restoration costs and benefits using an integrated framework that accounted for ecosystem services (e.g., carbon sequestration and water purification), as well as commercial and recreational fisheries across different reef patches. Assuming a one-time restoration intervention (i.e., urchin culling and kelp seeding), the authors found that the expected benefits justified restoration actions, while also acknowledging that the cost-benefit ratio substantially varies across patches (i.e., 0.33–3.4) depending on their ecological and logistical characteristics. Arroyo-Esquivel et al. (2023) developed the first spatial population biology model of kelp forests to evaluate the potential effectiveness of different restoration strategies, including purple urchin removal and external kelp outplanting, under a fixed level of ongoing intervention.

While Carnell et al. (2025) provides an economic justification for kelp forest restoration, the analysis is limited to a one-time intervention which does not address the reality that restoration will typically be a multi-year effort. Although Arroyo-Esquivel et al. (2023) offers a spatially explicit ecological model, the approach does not incorporate restoration costs, and the fixed restoration intensity within a framework does not address the reality of evolving restoration needs. Both studies assume deterministic restoration outcomes and do not address the strong stochasticity of the kelp system. While these analyses assess the efficiency of one or a limited set of management options, to the best of our knowledge, no existing study has solved for

optimal kelp forest restoration.

The restoration literature acknowledges that restoring ecosystems is expensive due to the labor-intensive nature of the work, need for monitoring and long-term maintenance, outcome uncertainty, and complexity of biological and environmental interactions within the habitat (Kimball et al., 2015; Bayraktarov et al., 2016; Brancalion et al., 2019). While restoration strategy should ideally consider expected costs and benefits, key species dynamics (including interactions), and environmental uncertainty, to our knowledge, there has been no such analysis for kelp forest restoration. To address this gap, we combine these elements into an integrated bioeconomic model for insights into efficient use of restoration actions—like kelp outplanting and purple urchin culling—to maximize net benefits from the system.

Such an Ecosystem-Based Management approach, capturing feedback between species, resource management and stakeholder well-being, is admittedly complex. However, a growing body of literature illustrates the importance of such a wide lens. For instance, Sanchirico and Springborn (2011) proposed a framework to evaluate the positive effects of mangrove habitats on fish stocks and ecosystem services, estimating the trade-offs between mangrove conservation and coastal development. Kahui et al. (2016) developed a bioeconomic model to examine the impact of destructive fishing gear on cold-water coral habitats. Similarly, Vondolia et al. (2020) constructed a dynamic bioeconomic model to analyze economically optimal harvest rules for a cod fishery and kelp forests, considering the ecosystem services of kelp forests that influence cod stock growth and environmental carrying capacity. These studies demonstrate that habitat conservation can enhance fishing revenue and the economic well-being of stakeholders. While such integrative approaches have been applied to various ecosystems, as summarized above, their application to kelp forest management is limited.

We assess kelp restoration strategies for balancing the long-run economic benefits to coastal communities with the cost of restoration investments under environmental uncertainty, particularly the future frequency and stochastic occurrence of MHWs—periods when sea surface temperatures exceed the local 90th percentile baseline for days to months—which simultaneously warm and de-nutlify coastal waters, causing catastrophic losses of canopy-forming kelps (McPherson et al., 2021). Also, available kelp biomass, as a food source for urchins, alters urchin foraging behavior (Karatayev et al., 2021; Arroyo-Esquivel et al., 2023), the initial temperature shock cascades into a nonlinear grazer–kelp feedback that can lock systems into degraded states. Models that ignore the stochastic and detrimental impacts of marine heatwaves (MHWs), the uncertainty surrounding future MHW regimes, and the unique behavioral dynamics of kelp–grazer systems may yield misleading restoration and policy recommendations. Explicitly incorporating MHW uncertainty, regime-shift risk, and ecosystem-specific dynamics is essential for designing adaptive management strategies for kelp forests increasingly exposed to climate volatility.

In this paper, we identify economically optimal kelp restoration strategies, which we reconcile short-term restoration costs with long-term economic gains. Our approach integrates three components: (1) the biology of kelp forests and key invertebrate species that depend on them, (2) the associated economic benefits and costs, and (3) the variable ocean climate of Northern California. We address three key questions: (i) Under realistic economic and ecological conditions, what is the optimal state-dependent restoration intensity in a system with non-linear kelp-urchin feedbacks? (ii) How do optimal restoration strategies adapt across different MHW frequency regimes? (iii) How does optimal restoration behavior shift under alternative ecological and economic scenarios? By identifying optimal balanced restoration activities for kelp forest restoration through a dynamic bioeconomic restoration model, this research can inform effective management strategies and support the long-term sustainability of this vital coastal ecosystem.

2 Bioeconomic model of population dynamics and optimal management in the kelp-urchin system

The dynamic bioeconomic model tracks the biomass of kelp (X_t , number of stipes), purple urchins (P_t , number of individuals), and red urchins (R_t , number of individuals) over annual time steps (t) within a representative restoration patch. The manager can influence the dynamics of these three state variables by selecting annual levels of two restoration actions: kelp outplanting intensity ($A_{X,t}$) and purple urchin removal effort ($A_{P,t}$). Both actions are measured in “labor hours,” encompassing both preparation and implementation tasks (e.g., preparation and lab culturing of juvenile kelp for outplanting and diving time).

As a starting point, we adapt the deterministic dynamic kelp-urchin structure of Arroyo-Esquivel et al. (2023). We extend the model in three ways. First we add the red urchin population in addition to the purple urchins originally modeled. Second we add uncertainty and learning to the model in the form of stochastic annual occurrence of MHWs, initial uncertainty about the true MHW regime (annual likelihood), and eventual identification of the true MHW regime. Finally, management actions in Arroyo-Esquivel et al. (2023) were limited to a set of static alternatives. Here we solve for the optimal feedback policy, i.e., the ideal levels of kelp outplanting and purple urchin removal depending on the current state of the system (given by the biological populations and current knowledge on the MHW regime). We develop the bioeconomic model for a representative patch of kelp forest.

2.1 Biological dynamics

A pivotal biological interaction in our model is the grazing of mature kelp by urchins. While our biological model broadly follows Arroyo-Esquivel et al. (2023), we adapt it to an annual time scale to reflect long-term

management planning rather than short-term recovery. Our annual time step starts in the fall season, aligning with the sequence of key ecological processes: kelp recruitment occurs during fall and winter, followed by natural kelp mortality in winter, urchin spawning in spring and summer, and restoration interventions in spring and summer (California Department of Fish and Wildlife, 2020; Ward et al., 2022). Also, the data used for model parameterization was primarily collected during the fall, further justifying this season as the start of our time step. The spatial unit for our model is a representative 360 m^2 patch, which matches the survey plot size in our parameterization dataset. We use the subscript t to indicate annual time steps for all state and control variables.

Three state variables indicate the levels of the key populations at the beginning of period t : X_t is the number of kelp stipes, P_t is the number of purple urchins, and R_t is the number of red urchins. The adult kelp remaining after grazing by urchins is:¹

$$K_t = k(X_t, P_t, R_t) = X_t \cdot \underbrace{\exp\left(-\gamma_P \cdot g(X_t, P_t) - \gamma_R \cdot R_t\right)}_{\text{grazing on adult kelp}}. \quad (1)$$

The coefficients γ_P and γ_R correspond to the rate of decline in kelp abundance per unit of grazing intensity from purple and red urchins, respectively. The function $g(\cdot)$, defining the grazing intensity from purple urchins, is described in detail next.

Under a kelp-abundant environment, purple urchins subsist on “drift kelp” blades that detach from extant kelp and settle on the sea floor (Harrold and Reed, 1985; Kriegisch et al., 2019). This “passive” grazing on drift kelp does not result in kelp mortality. In contrast, at low kelp density, as could occur under heat stress or poor nutrient availability, purple urchins switch from passive to “active” grazing, i.e., consuming the live kelp plant, including the stipe and holdfast leading to kelp mortality.² This unique grazing behavior (active at low kelp density, passive at high kelp density) is represented by unimodal Holling’s Type IV functional response in resource-consumer dynamics (Arroyo-Esquivel et al., 2023; Karatayev et al., 2021). Using the Type IV response specification from Cui et al. (2016), the grazing intensity of purple urchin is given by³:

$$g(X_t, P_t) = \frac{P_t}{\sigma_1 \cdot X_t^2 + \sigma_2 \cdot X_t + 1}, \quad (2)$$

where σ_1 drives the kelp density at which purple urchin grazing intensity reaches its maximum and σ_2

¹Arroyo-Esquivel et al. (2023) modeled the grazing effect on kelp density using a multiplicative form with a *max* operator, for example, $K_t = X_t \cdot \max[0, \gamma_P \cdot g(X_t, P_t)]$. In contrast, we adopt a smooth formulation of grazing pressure by applying an exponential function.

²Kelp stipes, similar to stems in land plants, offer structural support, linking leaf-like blades to the plant’s base and facilitating upward growth towards sunlight. The holdfast, resembling a root, anchors the kelp to the ocean floor, usually on rocky substrates, securing it against currents and waves.

³Cui et al. (2016) originally adopts a continuous-time formulation. We convert this functional response into an annual discrete-time specification for our model.

regulates the magnitude of this grazing intensity. This functional response represents a unimodal and concave grazing intensity that peaks at $X_{\max \text{ graze}} \equiv \frac{1}{\sqrt{\sigma_1}}$, the kelp density at which grazing per unit of kelp is maximized for a given level of purple urchins. In contrast, red urchins exhibit only weak swarm behavior, with a less pronounced intensity compared to purple urchins and a significantly smaller observed population size. For simplicity, we assume red urchin grazing follows and Type I functional response, i.e., linear prey consumption behavior.

We assume that the recruitment of the kelp population is governed by a saturating Beverton-Holt type recruitment, primarily driven by competition among the kelp spores for available resources (e.g., settlement space and light availability for photosynthesis). This type of density-dependent recruitment is typically suitable for species with a high fertility rate but a very low proportion of which become adults (Clark, 1976), such as bull kelp. Given kelp spore production of ν per unit of adult kelp, the resulting level of viable adult kelp recruits per unit of existing adult kelp (K_t) is:

$$y(X_t, P_t, R_t) = \frac{\nu \cdot \exp\left(-\mu_P \cdot g(X_t, P_t) - \mu_R \cdot R_t\right)}{1 + \phi \cdot K_t}, \quad (3)$$

where μ_P and μ_R determine the rates at which purple and red urchins consume kelp spores, respectively. The coefficient ϕ denotes the intensity of density-dependent recruitment. When the kelp density is sufficiently high, the total recruitment level across the stock (K_t) would approach $\frac{\nu}{\phi}$ in the absence of urchin spore consumption.

The dynamics of the kelp populations in a given year are:

$$X_{t+1} = x(X_t, P_t, R_t, A_{X,t} | Z_t) = \underbrace{\left[k(X_t, P_t, R_t) \cdot \left(y(X_t, P_t, R_t) + \delta_X \right) \right]}_{\text{kelp recruitment \& adult survivorship}} \cdot \underbrace{(1 - \xi \cdot Z_t)}_{\text{MHW shock}} + \underbrace{\eta_X \cdot A_{X,t}}_{\text{management: outplanting}}, \quad (4)$$

which combines the Type IV grazing response, kelp recruitment $i(X_t, P_t, R_t)$, the restoration intervention of (juvenile) kelp outplanting $A_{X,t}$, density-independent natural mortality of adult kelp and a stochastic MHW shock. δ_X is the coefficient of natural annual survival of adult kelp and η_X is the marginal effect of juvenile kelp outplanting.

Z_t is a binary random variable for the occurrence of a MHW, which leads to mortality of a fraction, ξ , of kelp.⁴ The key stochastic component in the biological dynamics is the annual occurrence of a marine heatwave (MHW), which adversely affects kelp density. We model the presence of a MHW in year t as a

⁴Arroyo-Esquivel et al. assume that kelp recruitment occurs before grazing. Given their model's monthly time steps, this ordering is of only minor consequence. Since our model is annual, we instead assume that recruitment occurs after grazing so that grazing of adult kelp influences the level of kelp recruitment into the next year.

Bernoulli random variable $Z_t \in \{0, 1\}$, where $Z_t = 1$ denotes the occurrence of a MHW, with probability α . In our framework, MHW events are conceptualized as integrated shocks to kelp dynamics that reflect both elevated sea temperatures and reduced nutrient availability. This formulation is supported by strong empirical correlations between MHWs and weakened coastal upwelling in this system (Brodeur et al., 2019; Varela et al., 2021; Izquierdo et al., 2022).

To build on this stochastic foundation, we incorporate an additional layer of climate uncertainty by accounting not only for annual MHW variability but also for long-run uncertainty in the underlying MHW regime—that is, the probability α of MHW occurrence each year. While most existing models treat climate drivers as fixed or known, we allow for the possibility that the true MHW regime is itself unknown and subject to discovery. This dual approach ensures that our insights into restoration strategies remain robust under both short-term variability and long-term climate uncertainty.

We consider two layers of uncertainty: MHWs are random occurrences (as described above) and the MHW frequency α is initially unknown with certainty. We assume that $\alpha = \alpha_i$ can take one of three values: $i \in \{l, m, h\}$, for low, medium, and high, respectively. The low case corresponds to MHW frequency as historically observed over the five decades ending in 2010. The high case corresponds to the elevated MHW frequency observed 2010-2020. Finally, the medium case is set to the average of the low and high rates. Initially the true level $\alpha = \alpha_{i^*}$ is unknown but the manager has beliefs about the likelihood of each possibility reflecting the true long run MHW rate, $f(i = i^*)$. For example, if the manager believes that each MHW regime has an equal likelihood of being the truth, then $f(i) = 1/3 \forall i$. We assume that scientists uncover and report the true regime $i^* \in \{l, m, h\}$ in T years (e.g., $T = 15$).

To model urchin reproduction via energy gained from kelp grazing, we distinguish between energy obtained from grazing on adult kelp (addressed here) and drift kelp (discussed in the following section). We assume that energy intake from microscopic spore kelp, as modeled in Equation (3), is negligible and thus does not contribute meaningfully to reproduction, following the approach of Arroyo-Esquivel et al. (2023). Accordingly, we focus on energy acquired through adult and drift kelp consumption as the primary sources of reproductive energy for purple urchins.

Purple urchins, due to their numerical dominance within kelp beds, are assumed to consume adult kelp before red urchins. This priority is based on observed habitat patterns in which red urchins typically inhabit the peripheries of kelp beds, where they forage on drift kelp that accumulates in those areas (California Department of Fish and Wildlife, 2020). As a result, purple urchins exert the initial grazing pressure on adult kelp. According to the adult kelp grazing model in Equation (1), they consume an amount of kelp equal to $X_t \cdot (1 - \exp(-\gamma_P \cdot g(X_t, P_t)))$. By applying the energy conversion efficiency ϵ_1 , we calculate the

resulting reproduction from this adult kelp consumption as:

$$p_X(X_t, P_t) = \epsilon_1 \cdot X_t \cdot (1 - \exp(-\gamma_P \cdot g(X_t, P_t))). \quad (5)$$

In modeling red urchin reproduction dynamics, we adopt a functional structure similar to that of purple urchins, but with a type I response. We posit that red urchins consume adult kelp after grazing by purple urchins, leading to the following formulation for red urchin reproduction through adult kelp consumption:

$$r_X(X_t, P_t, R_t) = \epsilon_2 \cdot X_t \cdot \exp(-\gamma_P \cdot g(X_t, P_t)) \cdot (1 - \exp(-\gamma_R \cdot R_t)), \quad (6)$$

where ϵ_2 represents the energy conversion factor for red urchins, quantifying the efficiency of reproductive energy production through the consumption of adult kelp.

Furthermore, we model the reproductive gains of urchins from drift kelp consumption as a function of kelp abundance and urchin biomass. Red urchins are assumed to graze first due to their superior competitive ability, followed by purple urchins. The reproductive output from drift kelp consumption by red and purple urchins is given by:

$$r_D(X_t, R_t) = \epsilon_4 \cdot X_t \cdot (1 - \exp(-\tau_R \cdot R_t)), \quad (7)$$

$$p_D(X_t, P_t, R_t) = \epsilon_3 \cdot X_t \cdot \exp(-\tau_R \cdot R_t) \cdot (1 - \exp(-\tau_P \cdot P_t)), \quad (8)$$

where ϵ_4 and ϵ_3 are effective energy conversion parameters for red and purple urchins, respectively, while τ_R and τ_P capture the strength of grazing effects by each urchin species on drift kelp. The full derivation of these equations and the biological assumptions underpinning them are provided in Appendix A.1.

The dynamics of purple urchins are then:

$$p(X_t, P_t, R_t, A_{P,t}) = \underbrace{\delta_P \cdot P_t}_{\text{urchin survival}} + \underbrace{p_X(X_t, P_t) + p_D(X_t, P_t, R_t)}_{\text{urchin recruitment}} - \underbrace{\eta_P \cdot A_{P,t}}_{\text{management: removal}}, \quad (9)$$

where δ_P is the annual survival proportion of purple urchins, and η_P quantifies the marginal effect of removal efforts on the purple urchin population (i.e., the number of purple urchins removed per diving hour). Restoration sites that undergo urchin removal often experience re-colonization due to exogenous urchin recruitment in subsequent periods (Ward et al., 2022). This dynamic reflects the tendency of patches with sufficiently low urchin density to rebound toward the densities of neighboring areas. To capture this

process, we model exogenous recruitment as occurring only when the urchin density in a given patch falls below a threshold. The resulting purple urchin population dynamics are expressed as follows:⁵

$$P_{t+1} = \max[P_{\text{exo}}, p(X_t, P_t, R_t, A_{P,t})], \quad (10)$$

where P_{exo} represents purple urchin exogenous recruitment for recolonization from adjacent patches.

We assume that harvest in the red urchin fishery is proportionally taken from the commercially valuable size class of the population, consistent with historical fishing practices. Specifically, red urchins are only harvested when they possess commercially valuable gonads, as red urchins in food-deprived environments lack sufficient gonadal quality for market purposes. Red urchin harvest is given by:

$$H_t = h(X_t, R_t) = \lambda \cdot \omega \cdot R_t \cdot \mathbb{1}(X_t > X_{\text{max graze}}), \quad (11)$$

where λ represents the proportion of the red urchin population that falls within the commercially harvestable size class and ω represents the annual fishing mortality rate. We assume that the kelp threshold above which red urchin gonads (or “uni”) attain commercial quality coincides with the kelp biomass level where grazing intensity per purple urchin is maximized, i.e., $X_{\text{max graze}}$, as specified in Equation (2). This is consistent with the observation that in purple urchin barrens—where kelp is low and direct grazing of kelp stipes is predominant—red urchins typically lack commercially viable gonads due to insufficient nourishment (Smith et al., 2021; Harrold and Reed, 1985).

Incorporating the previously established harvest rule and reproductive functions in Equations (6) and (7), along with natural mortality and recruitment, red urchin population dynamics are given by:

$$r(X_t, P_t, R_t) = \delta_R \cdot R_t + \left(r_X(X_t, P_t, R_t) + r_D(X_t, R_t) \right) - H_t, \quad (12)$$

where δ_R is the annual survival proportion of red urchins. Assuming that the red urchin population undergoes exogenous recruitment when its density falls below a certain threshold—similar to the dynamics of purple urchins—the final population dynamics of red urchins are given by:

$$R_{t+1} = \max[R_{\text{exo}}, r(X_t, P_t, R_t)], \quad (13)$$

⁵This formulation is supported by observations in the Noyo and Albion regions on the California coast, where annual monitoring of restoration sites revealed re-colonization and partial rebounds in urchin density following removal (Ward et al., 2022). In contrast, control sites showed no significant change in purple urchin density during the same period. Additionally, recolonization may be driven by food competition and self-avoidance behavior linked to disease outbreaks in densely populated areas (Dumont et al., 2007; Hereu et al., 2012). Consequently, we assume that when a given patch is sufficiently empty relative to neighboring patches, urchin density homogenizes to the level of adjacent patches.

where R_{exo} represents red urchin exogenous recruitment from neighboring patches.

2.2 Management components

Building on the biological dynamics and restoration management actions—kelp outplanting (A_X) and purple urchin removal (A_P)—described in the previous section, we now elaborate the cost structure of these restoration actions, as well as the benefits from red urchin harvest and the direct value of kelp.

We assume that the cost of implementing a restoration action j , for both kelp outplanting ($j = X$) and purple urchin removal ($j = P$), is linear:

$$c_j(A_{j,t}) = \theta_j \cdot A_{j,t}. \quad (14)$$

When red urchin harvest does take place, the cost is a fixed cost, independent of density and harvest level:

$$c_H(H_t) = \theta_H \cdot \mathbb{1}(H_t \neq 0). \quad (15)$$

This reflects the nature of diving-based harvest, where effort—such as descent, search, and handling—is roughly constant regardless of the quantity collected. Harvest occurs when two conditions are met: red urchins have commercially valuable gonads (as described in the previous section) and harvest revenue exceeds the harvesting cost, $\rho \cdot H_t > c_H$, where ρ is the ex-vessel price of red urchins.

On the benefit side of restoration, we consider two components. First, the red urchin fishery generates economic benefits not only from direct harvest revenue but also from broader economic multiplier effects in the local coastal economy through backward and forward linkages (e.g., seafood processing). Second, we assume that standing kelp biomass provides direct ecological and use value, modeled as a linear function $\kappa_X \cdot X_t$. The net benefit (reward) function in period t is:

$$\pi(H_t, X_t, A_{X,t}, A_{P,t}) = \underbrace{\kappa_R \cdot (\rho \cdot H_t - c_H(H_t))}_{\text{red urchin harvest profit and multiplier effects}} + \underbrace{\kappa_X \cdot X_t}_{\text{direct kelp value}} - \underbrace{c_X(A_{X,t})}_{\text{kelp outplanting cost}} - \underbrace{c_P(A_{P,t})}_{\text{purple urchin removal cost}}. \quad (16)$$

To calibrate the marginal direct kelp value (κ_X), we assume that the maximum direct kelp value is equal to the maximum achievable red urchin fishing benefits (profit and local economy spillover benefits). Further details are provided in Appendix C.3.

2.3 The management problem

Building on the biological dynamics and management components described above, the manager's objective is to choose the intensity of the two restoration actions in each period to maximize the expected present value of long-run net benefits:

$$\max_{A_{X,t}, A_{P,t}} \sum_{t=0}^{\infty} \beta^t \cdot E_Z [\pi(H_t, A_{P,t}, A_{X,t})] \quad (17)$$

$$s.t. \ X_{t+1} = x(X_t, P_t, R_t, A_{X,t} | Z_t)$$

$$P_{t+1} = p(X_t, P_t, R_t, A_{P,t})$$

$$R_{t+1} = r(X_t, P_t, R_t) \quad (18)$$

$$H_t = h(X_t, R_t)$$

$$Z_t \sim \text{Bernoulli}(\alpha_i)$$

$$Pr(i) = \begin{cases} f(i), & \text{if } t \in [1, T-1] \\ 1 \text{ for } i = i^* \text{ (given)}, & \text{if } t \geq T, \end{cases}$$

where β is a social discount factor, α is unknown for periods $t \in [0, T-1]$, and then α is known with certainty at the beginning of period T . We solve the problem using dynamic programming in two stages. First, for each possible regime $i \in \{l, m, h\}$, we solve for the optimal management policy over the known-MHW-regime stage, $t \in [T, \infty)$. Second, conditional on those solutions (and beliefs regarding the likelihood of each regime), we solve for the optimal management policy over the unknown-MHW-regime stage, $t \in [0, T-1]$.

2.3.1 Second stage: known-MHW-regime

The infinite-horizon Bellman equation for the known-MHW-regime stage is given by:

$$V(X_t, P_t, R_t | i^*) = \max_{A_{X,t}, A_{P,t}} \pi(H_t, A_{X,t}, A_{P,t}) + \beta E_Z [V(X_{t+1}, P_{t+1}, R_{t+1} | i^*)] \quad (19)$$

$$s.t. \ X_{t+1} = x(X_t, P_t, R_t, A_{X,t} | Z_t)$$

$$P_{t+1} = p(X_t, P_t, R_t, A_{P,t})$$

$$R_{t+1} = r(X_t, P_t, R_t) \quad (20)$$

$$H_t = h(X_t, R_t)$$

$$Z_t \sim \text{Bernoulli}(\alpha_{i^*}).$$

Because the MHW regime (i^*) is known with certainty, the expectation of the value function on the right hand side is given by:

$$\begin{aligned}
E_Z[V(X_{t+1}, P_{t+1}, R_{t+1}|i^*)] &= \sum_{Z_t \in \{0,1\}} Pr(Z_t|i^*) \cdot V(x(X_t, P_t, R_t, A_{X,t}|Z_t), p(X_t, P_t, R_t, A_{P,t}), r(X_t, P_t, R_t)) \quad (21) \\
&= \alpha_{i^*} \cdot V(x(X_t, P_t, R_t, A_{X,t}|Z_t = 1), p(X_t, P_t, R_t, A_{P,t}), r(X_t, P_t, R_t)) + \\
&\quad (1 - \alpha_{i^*}) \cdot V(x(X_t, P_t, R_t, A_{X,t}|Z_t = 0), p(X_t, P_t, R_t, A_{P,t}), r(X_t, P_t, R_t)).
\end{aligned}$$

2.3.2 First stage: unknown-MHW-regime

In the unknown-MHW-regime stage ($t \in [0, T-1]$) the manager does not know which regime i is the truth but instead has beliefs given by the density $f(i)$. Furthermore, we assume that during this stage the focus is on restoration and there is no harvest, $H_t = 0$ ⁶. Because the Bellman equation for the unknown-MHW-regime stage reflects a finite time horizon, the value function—in this stage given by J —depends on the period and is indexed by t :

$$J_t(X_t, P_t, R_t) = \max_{A_{P,t}, A_{X,t}} \pi(X_t, A_{X,t}, A_{P,t}) + \beta E_{Z,i}[J_{t+1}(X_{t+1}, P_{t+1}, R_{t+1})] \quad (22)$$

$$s.t. \ X_{t+1} = x(X_t, P_t, R_t, A_{X,t}|Z_t)$$

$$P_{t+1} = p(X_t, P_t, R_t, A_{P,t})$$

$$R_{t+1} = r(X_t, P_t, R_t) \quad (23)$$

$$H_t = 0$$

$$Z_t \sim \text{Bernoulli}(\alpha_i)$$

$$Pr(i) = f(i).$$

With regime uncertainty inherent in MHW, the expected value function on the RHS for all $t \in [0, T-1]$ is

$$\begin{aligned}
E_{Z,i}[J_{t+1}(X_{t+1}, P_{t+1}, R_{t+1})] &= \sum_{i \in \{l,m,h\}} Pr(i) \cdot E_Z[J_{t+1}(X_{t+1}, P_{t+1}, R_{t+1}|i)] \quad (24) \\
&= \sum_{i \in \{l,m,h\}} f(i) \left\{ \alpha_i \cdot J_{t+1}(x(X_t, P_t, R_t, A_{X,t}|Z_t = 1), p(X_t, P_t, R_t, A_{P,t}), r(X_t, P_t, R_t)) + \right. \\
&\quad \left. (1 - \alpha_i) \cdot J_{t+1}(x(X_t, P_t, R_t, A_{X,t}|Z_t = 0), p(X_t, P_t, R_t, A_{P,t}), r(X_t, P_t, R_t)) \right\}.
\end{aligned}$$

⁶Following the collapse of kelp forests and the subsequent fishing disaster declaration for the red urchin fishery in Northern California in 2016, there was not only a significant reduction in harvest but also an exodus of urchin divers. Red urchin landings and the number of divers have remained low relative to pre-MHW levels. Therefore, we assume that during the restoration period, actual harvest remains negligible due to the continued absence of returning divers and the lack of a sufficient food source.

For time $t = T$, facilitating the transition between the two stages, the expected value function is

$$E_{Z,i}[J_T(X_T, P_T, R_T)] = \sum_{i \in \{l, m, h\}} Pr(i) \cdot E_Z[V(X_T, P_T, R_T|i)], \quad (25)$$

where $E_Z[V(\cdot)]$ on the RHS is given by Equation (21).

2.4 Parameterization

We identified plausible parameter values for the model using peer-reviewed literature, government reports, personal communications with experts, and survey data from kelp forest systems. Details on the selection and estimation of all parameters are provided in Appendix C.

For estimating the biological parameters, we leveraged annual species density data collected through SCUBA surveys by Reef Check, a marine ecosystem conservation NGO. These surveys span the period prior to the intensification of marine heatwaves (2007–2012) and include six sites across Northern California, covering three focal species: kelp, purple urchin, and red urchin. Each observation in this dataset corresponds to a spatial unit of 360 m², comprising six adjacent 60 m² transects. For consistency, we parameterize our bioeconomic model at the scale of this representative patch (360 m²). To ensure that the biological parameters align with the observed species dynamics, we employ Approximate Bayesian Computation (ABC), as detailed in Appendix C.1.

Data sources and calculations for MHW frequencies, as well as parameters related to fishery and restoration management, are described in Appendices C.2 and C.3, respectively. Definitions of all state variables, along with parameter values and sources, are summarized in Table 1, encompassing biological parameters (estimated via ABC or drawn from the literature) and fishery- and economic-related parameters relevant to restoration.

Table 1: Variable and parameter definitions with values and sources

	Description	Value	Source
Variables			
t	Annual timestep		
X_t	Stock of adult kelp (# stipes) at time t		
P_t	Stock of purple urchins (# urchins) at time t		
R_t	Stock of red urchins (# urchins) at time t		
$A_{X,t}$	Hours of outplanting effort for X at time t		
$A_{P,t}$	Hours of removal effort targeting P at time t		
Z_t	Stochastic marine-heatwave (MHW) indicator		
Parameters			
γ_P	Per-unit grazing intensity of P on X	0.00036	ABC
γ_R	Per-unit grazing intensity of R on X	0.00039	ABC
σ_1	Peak location of grazing of P on X	5.0989e-05	ABC
σ_2	Magnitude of grazing of P on X	0.00505	ABC
μ_P	Per-unit grazing intensity of P on X spores	0.00041	ABC
μ_R	Per-unit grazing intensity of R on X spores	0.00090	ABC
ϕ	Beverton–Holt density-dependence for recruitment	0.00077	ABC
ϵ_1	Conversion factor for reproduction of P via grazing	2.2761	ABC
ϵ_2	Conversion factor for reproduction of R via grazing	0.8841	ABC
ϵ_3	Conversion $\times$ drift kelp reproduction-loss ratio (for P)	0.7230	ABC
ϵ_4	Conversion $\times$ drift kelp reproduction-loss ratio (for R)	1.2405	ABC
ν	Per-unit spore production of X	1.346	ABC
τ_P	Per-unit grazing intensity of P on drift X	0.00127	ABC
τ_R	Per-unit grazing intensity of R on drift X	0.00024	ABC
ξ	Fractional mortality of X during marine heatwave	0.4	Personal comm.
η_X	Marginal effect of $A_{X,t}$	2.13 kelps/hr	Carney et al. (2005)
η_P	Marginal effect of $A_{P,t}$	319.8 urchins/hr	Ward et al. (2022)
δ_X	Annual survival proportion of X	0.15	Personal comm.
δ_P	Annual survival proportion of P	0.88	Ebert (2010)
δ_R	Annual survival proportion of R	0.93	Ebert et al. (1999)
P_{exo}	Exogenous recolonization of P (per patch)	53 urchins/patch	Inferred from Ward et al. (2022)
R_{exo}	Exogenous recolonization of R (per patch)	25 urchins/patch	Assumed equal rate to P_{exo}
λ	Fraction of harvestable urchins	0.55	Morgan et al. (2000)
ω	Annual fishing mortality of R	0.4	Morgan et al. (2000)
ρ	Ex-vessel price per R	\$3.4/urchin	CDFW (2020)
θ_X	Unit cost of outplanting effort for X (\$/hr)	\$42.5/hr	Carney et al. (2005)
θ_P	Unit cost of removal effort for P (\$/hr)	\$102.60/hr	Ward et al. (2022)
θ_H	Fixed harvest cost of R (per patch)	\$16.21/patch	Personal comm. and Fishman (2022)
κ_R	Revenue multiplier from R fishery to community	\$3.1/dollar	Jacobsen et al. (2014)
κ_X	Economic value per unit of X	\$1.63/dollar	Set equal to commer. fish. value
β	Discount factor	0.97	Assumed
α_i	MHW occurrence probability in regime $i \in \{l, m, h\}$	{0.04,0.156,0.27}	Informed by McPherson et al. (2021); NOAA Physical Sciences Laboratory (2025)

3 Results

The solution to our dynamic programming problem is given by two functions. The policy function shows the optimal management action given a particular state. The value function shows, for any given state, the present expected value of long-run net benefits given optimal management. Because we have three state variables, it is not possible to visualize one of these functions in a single plot. Instead, we show three heatmap panels, where each state variable is held constant, in turn, while the other two vary along the two axes. For selecting the level of the state variable that is latent (held fixed) in each heatmap, we consider three alternatives. The first reflects a degraded ecosystem, roughly consistent with recent conditions in northern California following the decline of kelp. We specify these levels as a percentage of the maximum level we allow for each variable. This degraded scenario is given by kelp at 25%, purple urchins at 75% and red urchins at 25% of their maximum biomass. Defining script variables as indicating percentage of maximum biomass, this degraded scenario is given by $(X = 25, P = 75, R = 25)$. The second alternative we consider is a restored case given by the vector $(X = 75, P = 25, R = 75)$. Finally, we consider the midpoint between degraded and restored: $(X = 50, P = 50, R = 50)$.

In Table 2 we show the model components that allow for alternative specifications. We begin in the next section by presenting results for a “baseline scenario” as indicated by the underlined specification in the table. This scenario reflects an infinite time horizon perspective with MHW regime uncertainty resolved and red urchin harvesting allowed (i.e., the exploitation stage) solved at the central alternative for MHW regime and shown in figures at the central alternative for latent states. After establishing this baseline, we present results for alternative specifications focusing on results that differ from the baseline.

Table 2: Alternative model specifications

Component	Alternatives		
MHW regime	low	<u>medium</u>	high
MHW regime knowledge	uncertain	<u>certain</u>	
Red urchin harvest	not allowed	<u>allowed</u>	
Time horizon	finite	<u>infinite</u>	
External recruitment	available	<u>not available</u>	
Latent states in plotting*	degraded	<u>midpoint</u>	restored

Note: The “restoration stage” specification is shown in green; “exploitation stage” specification is shown in blue; and “baseline scenario” specification is underlined.
 *This option does not affect the model solution, but rather specifies the subset of the overall results shown in figures.

3.1 Baseline scenario

In Figure 1 we show the value function for our baseline scenario, i.e., for the exploitation stage under a known, medium MHW regime with red urchin harvest allowed and latent states fixed at their midpoint. As expected the value function is increasing in kelp biomass and decreasing with purple urchin biomass. There is a sharp jump as kelp increases past $X_{\max \text{ graze}}$ (approximately $X = 130$) the threshold past which we assume red urchin gonads and thus harvest become commercially viable (see Equation (11)). This threshold also marks the point of maximum purple urchin active grazing intensity (per unit of kelp biomass) as specified by the Holling Type IV functional response in Equation (2).




Figure 1: Value function for the baseline scenario (medium MHW regime; exploitation stage; midpoint latent states), shown as a heatmap for each pair of state variables on the vertical and horizontal axes.

In contrast to the other two stocks, the value function can be increasing or decreasing in red urchin biomass, depending on the state. This is not surprising given that red urchin biomass generates rewards from harvest and also losses from grazing valuable kelp. In the first subplot of Figure 1, we see that increasing red urchin biomass leads to decreasing value when kelp is low (e.g., $X < 150$). But when kelp is high this changes—value increases in red urchin biomass at first before falling. This shows that the beneficial role of additional red urchins dominates when their density is low to medium and the kelp population is robust. From the middle subplot we also see that the net beneficial role of red urchins attenuates as purple urchin biomass increases (and combined grazing pressure builds). This pattern highlights a trade-off between the commercial harvest value of red urchins versus kelp value, where the returns to kelp stem from both direct benefits of the kelp itself and indirect benefits from supporting the red urchin population.

Figure 2 shows the policy functions for the intensity of kelp outplanting (A_X , 2a) and the removal of purple urchins (A_P , 2b). The optimal outplanting function is generally intuitive: outplanting intensity is decreasing in kelp biomass and modestly increasing in the levels of both urchin species, which graze on kelp. outplanting increases consistently as kelp biomass declines below levels around $X_{\max \text{ graze}}$ (the level of kelp


at which grazing per unit of kelp is at its maximum for any given level of purple urchins). This transition happens at a higher kelp level when urchin populations are high and vice versa. Outside of this zone, at higher kelp levels, outplanting is generally at or near zero.

This shift—from intensive outplanting at low kelp biomass to near-zero effort at higher levels—stems from economic efficiency and kelp recovery dynamics. When kelp is scarce (especially near $X_{\max \text{ graze}}$), outplanting yields high marginal benefits by accelerating recovery and enabling both direct kelp value and red urchin fishery potential, particularly where urchins lack commercial value due to limited gonad development from food scarcity. Thus, outplanting in these states boosts long-run community benefits, as reflected by the value function’s sharp increase near $X_{\max \text{ graze}}$. In contrast, once kelp biomass surpasses a threshold, natural recruitment may sustain the system, and the marginal benefit of outplanting declines. Further intervention becomes inefficient, with costs outweighing returns—explaining the near-zero outplanting policy in high kelp states.


In contrast to outplanting, optimal purple urchin removal intensity (A_P)—shown in Figure 2b—is generally non-zero, except when purple urchin biomass (P) is very low. Optimal removal is also insensitive to red urchin levels, only mildly decreasing in kelp biomass, and, intuitively, increasing in purple urchin biomass. Counterintuitively, we find that as kelp biomass declines, the responsiveness to increasing purple urchin levels *softens* as seen in the final subplot of Figure 2b. To see why this is the case, note that the value function shown over P and X —last subplot of Figure 1—shows that the rate of change of the value with respect to P (vertical) is steeper when X is high. This shows that the marginal value of a decrease in purple urchins (e.g., from removal action) is greater in high-kelp versus low-kelp states. Since the cost of purple urchin removal does not depend on kelp biomass, this outcome is driven solely by net rewards from the biological implications of removals. Even though purple urchin grazing intensity per unit of purple urchin is falling as kelp increases past around $X_{\max \text{ graze}}$, the amount of kelp subject to grazing is increasing; due to this latter effect, as kelp increases the optimal purple urchin removal intensity does not attenuate with kelp recovery as we might expect, but rather stays constant or slightly increases.⁷

Figure 3 presents density functions for kelp and urchin biomass levels after 50 years, under no restoration (3a) or with optimal restoration (3b) given by the policy functions in Figure 2. Under no restoration, Figure 3a shows that both purple and red urchin populations could reach a wide range of levels but are most likely to be low. This outcome is mainly driven by scarcity of kelp biomass as a food source, but also the occasional MHW. Although kelp biomass also could vary widely, its density is bimodal with most likely levels at low and high extremes. While the univariate densities in Figure 3 are straightforward, they do not show how state

⁷In the adult kelp and purple urchin grazing dynamics—see Equations (1) and (2)—we specify that purple urchin grazing of kelp follows a Holling Type IV functional response, which results in a unimodal, concave grazing intensity over kelp biomass X , with peak grazing intensity per unit of purple urchin occurring at $X_{\max \text{ graze}} = \frac{1}{\sqrt{\sigma_1^2}}$.



(a) Kelp outplanting intensity, $A_X(X, R, P)$



(b) Purple urchin removal intensity, $A_P(X, R, P)$

Figure 2: Optimal restoration levels for the two management options—kelp outplanting (top row) and purple urchin removal (bottom row)—for the baseline scenario (medium MHW stage; exploitation stage; midpoint latent states).

levels are likely to co-vary across the three species. To show this, in Appendix Figures D14b (no restoration) and D15b (optimal restoration) we present bivariate densities in heatmap form. Under no restoration, Figure D14b shows (intuitively) that when the kelp population is high, urchin populations tend to be very low.

In contrast to the no restoration case, Figure 3b shows that under optimal restoration (outplanting and urchin removal) kelp and red urchin biomass are likely to be at higher levels, while purple urchins remain likely to be suppressed, in this case due to targeted removal instead of low kelp levels. In this ongoing restoration case, the bivariate heatmap densities in Appendix Figure D15b show that high kelp biomass commonly co-occurs with low purple urchin and high red urchin levels.

Next, we consider a set of alternatives to our baseline scenario. While we show full results for each (often in Appendix) we focus our attention in the main text on highlighting results that depart notably from the




Figure 3: Probability density functions for state variable levels after 50 years under the baseline scenario, with no restoration (top row) and with the optimal restoration policy (bottom row). Each density is generated by applying the Markov transition matrix over 50 periods and integrating over the other two state variables.

baseline established above.

3.2 Alternative scenario 1: Latent state slices in degraded and restored states

In our baseline scenario results, we showed two-dimensional plots selected from our three-dimensional value and policy functions, where each latent state (i.e., the state not shown on the vertical or horizontal axis in any bivariate plot) is set to its midpoint. Here we assess whether our qualitative results are sensitive to this choice by considering results based on alternative latent states: the first represents a set of “degraded” state levels ($\mathcal{X} = 25, \mathcal{P} = 75, \mathcal{R} = 25$) and the second a set of “restored” state levels ($\mathcal{X} = 75, \mathcal{P} = 25, \mathcal{R} = 75$). This involves assessing alternative two-dimensional slices of the three-dimensional value and policy functions. These results are presented for the value function in Appendix Figure D4 and for the policy functions in D8 and D9. They show that numerical magnitudes of these functions can shift intuitively, e.g., when the latent state of kelp is higher the value function shifts higher. But overall they confirm that our qualitative results are unchanged.

In the R – P state space under the degraded kelp state ($\mathcal{X} = 25$), outplanting appears where none was

observed in the midpoint slice. This pattern is consistent with the baseline policy function in Figure 2a, where outplanting intensity increases with purple and red urchin biomass when kelp is low ($X < 200$). Vertical slices at $X = 25$ in the X - R and X - P subfigures confirm that this behavior reflects underlying dynamics, not any new behavioral shift due to latent state choice. Overall, this indicates that our conclusions from the baseline model are not sensitive to the specific slice of the state space used for evaluation.

3.3 Alternative scenario 2: Restoration stage results

In our baseline model, we assume the MHW regime is known to be medium and red urchin harvest is allowed. Here we examine sensitivity of results to being in an initial restoration stage for $T = 15$ years in which red urchin harvest is not allowed and the MHW regime is unknown, assumed to be either low, medium or high with equal probability. We assume that at the end of this restoration stage, the true MHW regime becomes known (at one of the three possible levels), red urchin harvest resumes and management continues over an infinite planning horizon.⁸

When solving such a problem over a finite horizon (the initial restoration stage), the policy functions can vary over time (as the year approaches a point at which some components of the problem are expected to change, at time T). Even so, we find that these policy functions—shown at time $t = 1$ and $t = T$ in Appendix Figures D10 and D11—barely differ from our baseline solution. This is expected as t approaches T , since by $t = T + 1$ the model reverts to something very similar to the baseline model.⁹ The consistency from $t = 1$ through $t = T$ also indicates that the temporary moratorium on red urchin harvest does not have a strong effect on ideal restoration levels over that moratorium period. Still, we do observe that optimal outplanting and urchin removal intensities become slightly stronger over the restoration stage, showing that transitioning from expecting future harvest to realizing harvest incentivizes slightly stronger action.

To assess capacity for recovery over the restoration period, we calculated the probability density functions for state variable levels after $T = 15$ years, assuming that the system is at “degraded” state levels ($X = 25, P = 75, R = 25$) at time $t = 1$. We find that these densities—shown in Appendix Figure D18b—are visually almost indistinguishable from those of our baseline model after 50 years as shown in Figure 3b. This shows that even after a relatively short 15-year restoration period under the optimal policy, the system outcomes are very near the longer-run stationary distribution. The degraded kelp forest system modeled here can recover relatively rapidly once interventions begin.

⁸We solve this problem as follows. We first compute the value function solutions to the exploitation stage model (infinite horizon and red urchin harvest allowed) under each of the three MHW regime possibilities—see Section 2.3.1. We then combine these into a single value function by giving each an equal weight (1/3) and summing. This weighted combination then serves as the terminal value function for a finite horizon problem running from $t = 1$ to $t = T = 15$, which we solve with backwards induction—see Section 2.3.2.

⁹The only difference being that here we have equal weights across low, medium and high MHW regime levels versus simply the medium MHW regime in the baseline model.

3.4 Alternative scenario 3: Value and policy functions under different MHW regimes and regime beliefs

We examine how shifts in MHW frequency alter value and policy functions, shaping restoration patterns and outcomes. Relative to the baseline medium-MHW scenario, we evaluate how low and high MHW regimes influence optimal restoration, especially during the exploitation stage, to understand behavioral deviations under different climate conditions.

Appendix Figure D3 shows that value functions vary predictably across MHW regimes. Under the low-MHW scenario (Appendix Figure D3a), value increases across the state space due to reduced thermal stress, which promotes kelp growth and enhances ecosystem services and improved red urchin harvest. Conversely, the high-MHW regime (Appendix Figure D3b) yields lower values across all states, reflecting the compounded ecological stress of frequent MHWs.

One notable pattern under the high MHW scenario (Figure D3b) is the consistent decline in the value function with increasing red urchin biomass across all states. This contrasts with the low and medium MHW regimes, where the marginal net benefit of red urchins varies with kelp biomass—sometimes positive, sometimes negative. Under frequent MHW conditions, suppressed kelp growth diminishes the fishery-related benefit of additional red urchin biomass (as populations are more likely to fall below the $X_{\max \text{ graze}}$ threshold), making it insufficient to counterbalance the kelp losses from intensified grazing. Thus, red urchin biomass contributes positively to ecosystem value only when kelp biomass is high and stable—conditions more likely under less frequent MHW impacts.

The overall structure of the policy functions for both outplanting and purple urchin removal remains broadly consistent across MHW regimes. As shown in Appendix Figure D6, outplanting is concentrated in low kelp biomass states under all regimes, similar to the baseline medium-MHW case. However, regime-specific differences emerge in both intensity and spatial extent. Under low MHW (Appendix Figure D6a), reduced kelp mortality increases the marginal benefit of restoration, leading to more intensive outplanting—a pattern consistent with long-term investment logic under favorable conditions.

In contrast, under high MHW (Appendix Figure D6b), while peak outplanting intensity remains similar, restoration extends to new regions (e.g., around $X=250$ in the $X-R$ slice) that saw no intervention under milder regimes. This expanded effort, particularly when red urchin biomass is high, functions as a buffer strategy to prevent overgrazing from driving kelp biomass below the $X_{\max \text{ graze}}$ threshold, beyond which red urchin gonad quality declines and the fishery becomes economically nonviable.

Appendix Figure D7 shows that purple urchin removal is applied gradually across the entire state space, consistent with the baseline scenario. Compared to outplanting, removal intensity is relatively insensitive

to the MHW regime, showing minimal spatial variation. This likely reflects naturally suppressed urchin growth under intensified MHW due to limited kelp availability. Supporting this, long-run state distributions without restoration—Figure 3a (medium MHW) and Appendix Figure D16 (low and high MHW)—illustrate biomass dynamics across regimes. In particular, the low-MHW case (Appendix Figure D14a), which reflects pre-MHW conditions in Northern California, shows relatively stable coexistence of kelp, purple, and red urchins. However, under medium and high MHW, both urchin species increasingly struggle to persist without restoration.

However, as MHW severity increases under the medium and high MHW scenarios (Figures 3a and Appendix Figure D16a), kelp biomass generally shifts toward zero, and both urchin populations are severely degraded simultaneously. This pattern suggests that urchin biomass is strongly dependent on kelp availability, and that the absence of food naturally regulates urchin populations without the need for additional removal efforts. Consequently, purple urchin removal becomes economically inefficient under severely suppressed kelp conditions, where urchin reproduction is already limited.

Lastly, restoration managers may form heterogeneous or biased beliefs about the prevailing MHW regime, shaped by scientific forecasts or subjective expectations. In addition to the equal-belief case, we examine two alternative belief structures: (1) a low-MHW-biased belief, $\phi(i) = \{\frac{2}{3}, \frac{1}{6}, \frac{1}{6}\}$, and (2) a high-MHW-biased belief, $\phi(i) = \{\frac{1}{6}, \frac{1}{6}, \frac{2}{3}\}$, where $i = l, m, h$ represents the MHW regimes.


As expected, the resulting outplanting policy functions (Appendix Figure D12) closely mirror those under the corresponding regime assumptions, indicating that managers' beliefs guide restoration behaviors toward the anticipated long-run policy. In contrast, purple urchin removal remains largely unchanged—reinforcing its insensitivity to regime expectations. Final state distributions under biased beliefs (Appendix Figure D18) lie between those of the fully certain belief cases, as theoretically expected.

3.5 Alternative scenario 4: Restricted restoration


In practice, management agencies involved in restoration pilot programs or experiments often apply only one type of restoration intervention—either kelp outplanting or purple urchin removal. Even when both methods are deployed, they are typically implemented intermittently rather than continuously. Our baseline scenario demonstrated that, depending on the state of the kelp forest, simultaneous implementation of both strategies is often required to maximize long-run net benefits to coastal communities. However, the potential consequences of restricted or partial restoration, where only one method by either kelp outplanting or purple urchin removal, is applied—remain unclear, particularly in terms of the expected economic benefits to coastal economies.

We further investigate how net benefits to coastal communities and restoration patterns vary under restricted restoration scenarios, where only one type of intervention—either kelp outplanting or purple urchin removal—is implemented. Figure 4 presents the percent change in value functions between the unrestricted case (with both restoration efforts) and each restricted case, defined as $100\% \times (V^{\text{rest.}} - V^{\text{unrest.}})/V^{\text{unrest.}}$, where $V^{\text{rest.}}$ is the value function under restricted restoration and $V^{\text{unrest.}}$ is from the unrestricted case.

Overall, we find that limiting restoration to a single intervention method reduces the net benefits to coastal communities, as reflected in the loss of value function.



(a) Only kelp outplanting (A_X) is implemented; urchin removal is restricted ($A_P = 0$)



(b) Only purple urchin removal (A_P) is implemented; kelp outplanting is restricted ($A_X = 0$)

Figure 4: Percent decrease in the value function from restricting implementation of restoration to a single action, either outplanting only (a) or purple urchin removal only (b), relative to the unrestricted case (medium MHW regime; exploitation stage; midpoint latent states).

One notable observation in the percent change of the value function is the asymmetric decline associated with restricted restoration methods. When only kelp outplanting is applied (i.e., $A_P = 0$), as shown in Figure 4a, we observe a broader and more substantial decrease in the value function across the kelp and urchin state space. This loss is particularly pronounced in regions with higher purple urchin biomass. Red urchins appear

to have a more neutral effect; however, their contribution to value loss increases when purple urchin levels are also high (second plot).


An additional interesting pattern appears in the third subplot, which holds red urchin biomass fixed at $\mathcal{R} = 50$. Here, the absence of purple urchin removal leads to increased value loss as kelp biomass rises, especially under high purple urchin densities. Although initially counterintuitive, this is consistent with the biological dynamics of the system: high kelp biomass provides ample food, promoting population growth for both purple and red urchins. As shown in the baseline and alternative MHW scenarios, elevated purple urchin populations consistently reduce the value function, and red urchins may also do so under high densities occasionally. Without concurrent removal efforts to curb their growth—particularly for purple urchins—high kelp biomass may fuel reproduction, thereby reducing overall value with the consumed kelp and increased urchin biomass. These findings underscore that the marginal net benefit of kelp outplanting increases significantly when complemented by purple urchin removal, which helps regulate potential urchin population booms.

In contrast, when only purple urchin removal is implemented, the overall decline in the value function is less visually pronounced across most of the state space. However, we observe a steeper drop in value relative to the case where only kelp outplanting is applied—particularly in regions with low kelp biomass (X). This indicates that the absence of outplanting can negatively impact net benefits to coastal communities, though the loss is less severe unless urchin populations—both purple and red—remain highly persistent. These results suggest that while the simultaneous implementation of both interventions is ideal, the role of outplanting becomes especially critical when kelp biomass is low. This finding is consistent with the baseline scenario’s outplanting policy function, which showed increased responsiveness under degraded kelp conditions.


Figure 5 presents the policy functions under restricted restoration scenarios: kelp outplanting only (top two rows) and purple urchin removal only (bottom row). In Figure 5a, the kelp outplanting intensity shifts noticeably across the state space compared to the baseline scenario, showing substantial restoration effort in certain regions. By contrast, when only purple urchin removal is implemented, the changes in policy intensity are relatively minor, indicating that kelp outplanting tends to drive more significant shifts in restoration behavior.

The shift in the outplanting policy function under restricted restoration is explained by the behavior in value function in Figure D5a, where A_X is active but $A_P = 0$. The first subplot shows that the marginal cost of kelp grazing increases when both purple and red urchin biomass are high, as their combined grazing suppresses kelp value more than either species alone. As a result, value increases more slowly with kelp biomass in high-urchin states.

Compared to the baseline (Figure 2a), outplanting no longer occurs in state areas where both urchin species are abundant, especially in the $X-R$ and $X-P$ slices. This indicates that without purple urchin



(a) Kelp outplanting intensity, $A_X(X, R, P)$ under $A_P = 0$



(b) Purple urchin removal intensity, $A_P(X, R, P)$ under $A_X = 0$

Figure 5: Optimal restoration levels under restricted scenarios—only kelp outplanting ($A_P = 0$, top row) and only urchin removal ($A_X = 0$, bottom row)—for the baseline scenario (medium MHW regime; exploitation stage; midpoint latent states)

removal, outplanting becomes economically inefficient under excessive grazing pressure. Notably, the P – R slice shows that when red urchin biomass is low, outplanting may still be economically viable even with very high purple urchin densities, as long as moderate kelp biomass is present ($X = 50$).

Figure 5b shows that purple urchin removal under the restricted restoration case (i.e., $A_X = 0$) follows a similar pattern to the baseline scenario, with gradual application across most states. However, a key difference emerges: there is a small but visible no-restoration zone when kelp biomass is near zero, which does not appear in the baseline case. This can be explained by the kelp dynamics equation (4): without outplanting ($A_X = 0$), there is no exogenous source of kelp recruitment. Once the system is absorbed into a near-zero kelp state, kelp recovery on its own becomes infeasible. As such, urchin removal in these states offers no economic benefit, rendering it inefficient in the absence of kelp recovery potential.




Figure 6: Value function in the alternative case where red urchin harvest is not available (medium MHW regime; exploitation stage with no harvest; midpoint latent states), shown as a heatmap for each pair of state variables on the vertical and horizontal axes.

In conclusion, implementing only one type of restoration not only reduces the net benefits from restoration but also alters the optimal restoration pattern in the absence of the other. This highlights that the two restoration efforts are endogenous to one another—their optimal implementation is jointly determined, and each complements the effectiveness of the other.

3.6 Alternative scenario 5: No return of the red urchin harvest even after restoration

Even in the long run after restoration, red urchin harvests may not fully recover during the exploitation stage due to diver exit, lack of return, and the deterioration of physical and intellectual fishing assets (e.g., boats and knowledge) resulting from prolonged inactivity in the fishery. Considering that fishery benefits from red urchins may remain limited even after restoration, we solve an alternative scenario assuming no harvest benefit during the exploitation stage, unlike the baseline case where such benefits are allowed.

The resulting value function, based solely on the direct value of kelp, exhibits a different pattern with respect to red urchin biomass. Figure 6 shows the value function under this no-harvest scenario. As in the baseline case, the value function increases monotonically with kelp biomass and decreases with purple urchin biomass. However, unlike the baseline, it now decreases monotonically with red urchin biomass—reflecting the fact that red urchins impose only grazing pressure and no longer contribute positively to restoration benefits. While the overall scale of the value function declines, the decrease is modest. This suggests that, in our model, the direct value of kelp is the primary driver of long-run restoration benefits.

Figure D13 shows the policy function under the no-harvest scenario. For outplanting intensity, we observe a slight reduction but no substantial change in the optimal strategy. Similarly, purple urchin removal intensity

shows a minor contraction but remains largely unchanged. Given the dominant contribution of direct kelp value to net benefits, this minor change in the restoration pattern without red urchin harvest is expected. Under our model assumptions, restoration remains economically justified even when based solely on the intrinsic value of kelp.

3.7 Alternative scenario 6: Restoration with exogenous kelp spore recruitment

In addition to the internal recruitment dynamics characterized in the baseline scenario, we also consider the possibility of external recruitment from kelp refugia—such as patches near river mouths—where kelp persists due to reduced urchin biomass driven by lower salinity. To incorporate this, we introduce an external spore recruitment parameter, η_E , into the kelp dynamics equation. This parameter represents a constant influx of spores from external sources and is assumed to be subject to grazing by urchins in the same manner as internally-produced spores. We analytically derive a candidate value for η_E based on observations from a restoration case study in Noyo, Northern California. The modified kelp dynamics and details of this derivation are provided in Appendix A.2.

Figure 7 represents the value function with external spore recruitment. With external spore recruitment from outside the patch, this alternative value function not only increases in numerical magnitude relative to the baseline but also now increases monotonically across red urchin biomass, which differs from the baseline scenario.






Figure 7: Value function in the alternative case with external recruitment (medium MHW regime; exploitation stage; midpoint latent states), shown as a heatmap for each pair of state variables on the vertical and horizontal axes.

The change in the value function can be understood by examining the corresponding policy function changes under external spore recruitment. While the purple urchin removal policy in Figure 8b shows little change, the outplanting policy in Figure 8a exhibits a notable shift—we observe almost no kelp outplanting

in areas where it was previously applied extensively.¹⁰ With the availability of external spore recruitment, reliance on kelp outplanting is greatly reduced compared to the baseline scenario, where internal recruitment is the sole source. This reduction renders outplanting investment potentially inefficient. With minimal outplanting and its associated costs, the long-run net benefit from kelp forest restoration improves substantially, as reflected in the value function. Moreover, consistent external spore recruitment can enhance the relative economic contribution of the red urchin fishery to the long-run net benefits—external recruitment alleviates the trade-off between the direct value of kelp and the fishery benefits resulting from increased red urchin biomass.



(a) Kelp outplanting intensity under external recruitment



(b) Purple urchin removal intensity under external recruitment

Figure 8: Optimal restoration levels for the two management options with external recruitment, kelp outplanting (top row) and purple urchin removal (bottom row) (medium MHW regime; exploitation stage; midpoint latent states)

The long-run probability density function in Figure 9, compared to the baseline optimal restoration outcome in Figure 3b, shows somewhat similar results to the baseline restoration with intensive outplanting.

¹⁰In policy function slices at degraded and restored stages, when kelp biomass is low (i.e., X), there is a slight rebound in outplanting intensity; however, overall efforts remain near zero.

Relative to the baseline scenario, the long-run kelp biomass appears more evenly distributed with increased probability of lower kelp outcomes. Management primarily through purple urchin removal as "gardening" would be sufficient intervention to recover the degraded kelp forest under consistent external spore recruitment from neighboring patches.




Figure 9: Probability density functions for state variable levels after 50 years under the baseline scenario with external spore recruitment along with restoration. Each density is generated by applying the Markov transition matrix over 50 periods and integrating over the other two state variables.

4 Discussion

This study offers the first dynamic, feedback-control evaluation of efficient restoration and long-run management of the kelp–urchin system. We integrate ecological interactions, restoration costs, fishery revenues, socio-economic benefits, and marine heatwave (MHW) uncertainty into a unified decision framework. In the baseline result, the welfare-maximizing strategy features an intensive pulse of kelp outplanting in the low kelp state, paired with gradually applied purple urchin removal efforts across purple urchin biomass. Once kelp biomass exceeds a threshold (e.g., up to 40% of the maximum kelp biomass level in our model), further outplanting becomes economically inefficient. At the same time, continued grazer control remains efficient due to increasing grazing pressure with kelp abundance. This management rule holds across the other MHW regime scenarios, indicating that a climate-robust strategy is to apply intensive outplanting, particularly in low kelp states such as urchin barrens common along the Northern California coast, along with simultaneously applied gradual purple urchin removal to maintain low urchin densities. By translating complex dynamics into an optimal restoration rule corresponding to given states, the model provides a practical, economically effective, and climate-robust guideline for restoration planning.

Without intervention, we find that our modeled kelp–urchin system under the medium MHW regime (e.g., increased MHW frequency) predicts that kelp biomass, while still capable of reaching a wide range of levels, generally exhibits two polarized long-run outcomes: either persistently high or severely suppressed states. The biomass of grazers—both purple and red urchins—is highly likely to remain degraded, and this predicted long-run outcome becomes more pronounced under the high MHW scenario. In particular, the high likelihood of low red urchin biomass under our model implies that the commercial red urchin fishery will continue to struggle if the MHW regime intensifies relative to the historical levels of MHW and restoration efforts are absent.

In contrast, implementing the optimal restoration strategy, as predicted by our model, reverses this projected outcome (and/or the state of urchin barrens observed in the real world) of the kelp–urchin system, increasing the likelihood of more desirable long-run socio-ecological outcomes—namely, higher kelp biomass, lower purple urchin biomass, and higher red urchin biomass. The findings from our model underscore the importance of intervention: expected long-run net benefits—characterized by fishing revenue in the red urchin fishery, its spillover economic benefits to the local coastal economy, and the ecological and use value of kelp—are more likely to be supported and sustained through active restoration efforts.

Uncertainty about the future of the red-urchin fishery need not delay restoration action. Even if urchin divers exit permanently and local fishing assets or knowledge are lost, our analysis indicate that restoration remains welfare-positive under this pessimistic scenario. In our analysis, even eschewing red urchin fishery

rents and relying solely on the direct value of kelp¹¹, the optimal restoration strategy still yields higher net returns than inaction. Considering direct kelp value itself, such as non-market ecosystem services or broader community impacts, our results suggest that proactive restoration is expected to provide long-term economic and ecological gains.

Once restoration begins, a site initially in a barren state is expected to reach a stationary distribution with higher average kelp biomass within the restoration stage—approximately fifteen years ($t = 1 - 15$) in our model. This recovery timeline offers a useful benchmark for managers and coastal stakeholders to anticipate when visible ecological gains may emerge. It provides critical insights for related policy decisions under uncertainty, such as those concerning the red urchin fishery—recently declared a disaster—and the recreational red abalone fishery, which remains closed. For instance, understanding how quickly kelp forests recover following intervention may inform policymakers to consider the timing of reopening these fisheries. In this sense, our model offers a valuable foundation for designing follow-up management strategies tied to kelp forest restoration, both for policymakers and for stakeholders awaiting the ecosystem’s return.

Restoration is most efficient when kelp outplanting and purple urchin removal are both available; excluding either intervention significantly reduces expected net benefits relative to the dual-action benchmark. The lost value from relying on one intervention is asymmetric: using outplanting alone results in a substantial loss of value (around half) because planted kelp is rapidly consumed before successful recruitment in grazer-rich conditions. By contrast, relying on urchin removal alone produces more modest losses overall. However, when natural recruitment is especially limited, recovery with culling alone is very slow or even infeasible.

Finally, our results show that, when consistent external spore recruitment is available, long-run net benefits increase and reliance on kelp outplanting declines sharply, with purple urchin removal alone often sufficient to achieve recovery efficiently. While we do not explicitly model spatial dynamics, these results do suggest a rationale for focusing recovery efforts on grazer control in areas near kelp refugia where kelp spore-generation persists.

While our analysis considers several alternative scenarios—restricting restoration options, exclusion of prospective fishery benefits, and pessimistic marine heatwave (MHW) conditions—several limitations remain. First, spatial spillovers, e.g., through kelp spore recruitment and urchin recolonization, are treated exogenously rather than modeled through an explicit spatio-temporal dynamic framework (e.g., meta-population models). Second, the MHW-induced kelp mortality parameter, specifying the intensity of our key system stressor, is based on laboratory estimates and has not yet been validated in the field. Third, our model does not account for potential endogenous behavioral responses by resource users, such as urchin divers, even though restoration-induced changes in kelp forest conditions over time may influence their long-term

¹¹The direct value of kelp was assumed to be roughly on par with red urchin fishery profits (see Appendix C.3).

participation decisions. For instance, sustained kelp degradation could lead to red urchin population decline, prompting divers to exit the fishery permanently. Moreover, when divers leave the fishery, their fishing knowledge and physical assets—such as boats and gear—may deteriorate over time. This gradual loss of human and physical capital reduces the likelihood of re-entry, even if kelp and urchin populations later recover. Future extensions that incorporate explicit spatial patch dynamics, adaptive behavioral responses by fishers across restoration phases, and field-based validation of MHW effects are promising directions for advancing our understanding of effective kelp forest restoration strategies.

References

- A. A. Alahmadi, J. A. Flegg, D. G. Cochrane, C. C. Drovandi, and J. M. Keith. A comparison of approximate versus exact techniques for bayesian parameter inference in nonlinear ordinary differential equation models. *Royal Society Open Science*, 7, 3 2020. ISSN 20545703. doi: 10.1098/rsos.191315.
- J. Arroyo-Esquivel, M. L. Baskett, M. McPherson, and A. Hastings. How far to build it before they come? analyzing the use of the field of dreams hypothesis in bull kelp restoration. *Ecological Applications*, 33(4): e2850, 4 2023. ISSN 1939-5582. doi: 10.1002/EAP.2850.
- E. Bayraktarov, M. I. Saunders, S. Abdullah, M. Mills, J. Beher, H. P. Possingham, P. J. Mumby, and C. E. Lovelock. The cost and feasibility of marine coastal restoration. *Ecological Applications*, 26(4):1055–1074, 6 2016. ISSN 1051-0761. doi: 10.1890/15-1077.
- M. A. Beaumont. Approximate bayesian computation. *Annual Review of Statistics and Its Application*, 6: 379–403, 3 2019. ISSN 2326831X. doi: 10.1146/ANNUREV-STATISTICS-030718-105212.
- S. Bennett, T. Wernberg, S. D. Connell, A. J. Hobday, C. R. Johnson, and E. S. Poloczanska. The ‘Great Southern Reef’: Social, ecological and economic value of Australia’s neglected kelp forests. In *Marine and Freshwater Research*, volume 67, pages 47–56. CSIRO, 2016. doi: 10.1071/MF15232.
- P. H. Brancalion, P. Meli, J. R. Tymus, F. E. Lenti, R. M. Benini, A. P. M. Silva, I. Isernhagen, and K. D. Holl. What makes ecosystem restoration expensive? a systematic cost assessment of projects in brazil. *Biological Conservation*, 240, 12 2019. ISSN 00063207. doi: 10.1016/j.biocon.2019.108274.
- R. D. Brodeur, T. D. Auth, and A. J. Phillips. Major shifts in pelagic micronekton and macrozooplankton community structure in an upwelling ecosystem related to an unprecedented marine heatwave. *Frontiers in Marine Science*, 6, 2019. ISSN 22967745. doi: 10.3389/fmars.2019.00212.
- California Department of Fish and Wildlife. Red sea urchin, *Mesocentrotus franciscanus*, enhanced status report, 2020. URL <https://marinespecies.wildlife.ca.gov/red-sea-urchin/the-species/>.
- California Ocean Protection Council. *Interim Action Plan for Protecting and Restoring California’s Kelp Forests*. California Ocean Protection Council, 2 2021.
- P. E. Carnell, K. Whiteoak, M. Young, K. Critchell, S. Swearer, P. I. Macreadie, J. McIntyre, and E. A. Trembl. Prioritising investment in kelp forest restoration: A spatially explicit benefit-cost analysis in southern Australia. *Ecosystem Services*, 74, 8 2025. ISSN 22120416. doi: 10.1016/j.ecoser.2025.101739.

- L. T. Carney, J. R. Waaland, T. Klinger, and K. Ewing. Restoration of the bull kelp *nereocystis luetkeana* in nearshore rocky habitats. *Marine Ecology Progress Series*, 302:49–61, 11 2005. ISSN 01718630. doi: 10.3354/MEPS302049.
- C. W. Clark. *Mathematical bioeconomics: The optimal management of renewable resources*. John Wiley & Sons, Inc., New York, 1976. ISBN 0-471-15856-9. A Wiley-Interscience Publication, First edition.
- K. Csilléry, M. G. Blum, O. E. Gaggiotti, and O. François. Approximate bayesian computation (abc) in practice. *Trends in Ecology & Evolution*, 25:410–418, 7 2010. ISSN 0169-5347. doi: 10.1016/J.TREE.2010.04.001.
- Q. Cui, Q. Zhang, Z. Qiu, and Z. Hu. Complex dynamics of a discrete-time predator-prey system with Holling IV functional response. *Chaos, Solitons and Fractals*, 87:158–171, 6 2016. ISSN 09600779. doi: 10.1016/j.chaos.2016.04.002.
- D. Duggins, J. E. Eckman, C. E. Siddon, and T. Klinger. Interactive roles of mesograzers and current flow in survival of kelps. *Marine Ecology Progress Series*, 223:143–155, 11 . ISSN 0171-8630. doi: 10.3354/MEPS223143.
- C. P. Dumont, J. H. Himmelman, and S. M. Robinson. Random movement pattern of the sea urchin *strongylocentrotus droebachiensis*. *Journal of Experimental Marine Biology and Ecology*, 340:80–89, 1 2007. ISSN 00220981. doi: 10.1016/j.jembe.2006.08.013.
- T. A. Ebert. Demographic patterns of the purple sea urchin *strongylocentrotus purpuratus* along a latitudinal gradient, 1985-1987. *Marine Ecology Progress Series*, 406:105–120, 5 2010. ISSN 01718630. doi: 10.3354/meps08547.
- T. A. Ebert, J. D. Dixon, S. C. Schroeter, P. E. Kalvass, N. T. Richmond, W. A. Bradbury, and D. A. Woodby. Growth and mortality of red sea urchins *strongylocentrotus franciscanus* across a latitudinal gradient. *Marine Ecology Progress Series*, 190:189–209, 1999.
- A. M. Eger, E. Marzinelli, P. Gribben, C. R. Johnson, C. Layton, P. D. Steinberg, G. Wood, B. R. Silliman, and A. Vergés. Playing to the positives: Using synergies to enhance kelp forest restoration. *Frontiers in Marine Science*, 7, 7 2020. ISSN 22967745. doi: 10.3389/fmars.2020.00544.
- P. L. Fackler. MDPSOLVE: A MATLAB toolbox for solving Markov decision problems with dynamic programming user’s guide, 2011. URL <https://sites.google.com/site/mdpsolve/users-guide>. Accessed online.

- K. Fishman. Red urchin fishery still struggling on North Coast after kelp collapse; ‘everybody’s scared’, Aug. 2022. URL <https://mendovoice.com/2022/08/red-urchin-fishery-still-struggling-in-north-coast-after-kelp-collapse-everybodys-scared/>.
- S. Fredriksen, K. Filbee-Dexter, K. M. Norderhaug, H. Steen, T. Bodvin, M. A. Coleman, F. Moy, and T. Wernberg. Green gravel: a novel restoration tool to combat kelp forest decline. *Scientific Reports*, 10, 12 2020. ISSN 20452322. doi: 10.1038/s41598-020-60553-x.
- C. Harrold and D. C. Reed. Food availability, sea urchin grazing, and kelp forest community structure. *Ecology*, 66:1160–1169, 8 1985. ISSN 1939-9170. doi: 10.2307/1939168.
- C. D. Harvell, D. Montecino-Latorre, J. M. Caldwell, J. M. Burt, K. Bosley, A. Keller, S. F. Heron, A. K. Salomon, L. Lee, O. Pontier, C. Pattengill-Semmens, and J. K. Gaydos. Disease epidemic and a marine heat wave are associated with the continental-scale collapse of a pivotal predator (*Pycnopodia helianthoides*). *Science Advances*, 5, 1 2019. ISSN 23752548. doi: 10.1126/sciadv.aau7042.
- B. Hereu, C. Linares, E. Sala, J. Garrabou, A. Garcia-Rubies, D. Diaz, and M. Zabala. Multiple processes regulate long-term population dynamics of sea urchins on mediterranean rocky reefs. *PLoS ONE*, 7, 5 2012. ISSN 19326203. doi: 10.1371/journal.pone.0036901.
- E. J. Hind. A review of the past, the present, and the future of fishers’ knowledge research: A challenge to established fisheries science. *ICES Journal of Marine Science*, 72:341–358, 7 2014. ISSN 10959289. doi: 10.1093/icesjms/fsu169.
- A. J. Hobday, L. V. Alexander, S. E. Perkins, D. A. Smale, S. C. Straub, E. C. Oliver, J. A. Benthuyssen, M. T. Burrows, M. G. Donat, M. Feng, N. J. Holbrook, P. J. Moore, H. A. Scannell, A. S. Gupta, and T. Wernberg. A hierarchical approach to defining marine heatwaves. *Progress in Oceanography*, 141: 227–238, 2 2016. ISSN 00796611. doi: 10.1016/j.pocean.2015.12.014.
- S. Hynes, W. Chen, K. Vondolia, C. Armstrong, and E. O’Connor. Valuing the ecosystem service benefits from kelp forest restoration: A choice experiment from Norway. *Ecological Economics*, 179, 1 2021. ISSN 09218009. doi: 10.1016/j.ecolecon.2020.106833.
- P. Izquierdo, F. G. Taboada, R. González-Gil, J. Arrontes, and J. M. Rico. Alongshore upwelling modulates the intensity of marine heatwaves in a temperate coastal sea. *Science of the Total Environment*, 835, 8 2022. ISSN 18791026. doi: 10.1016/j.scitotenv.2022.155478.

- F. Jabot, T. Faure, and N. Dumoulin. EasyABC: Performing efficient approximate Bayesian computation sampling schemes using R. *Methods in Ecology and Evolution*, 4:684–687, 7 2013. ISSN 2041210X. doi: 10.1111/2041-210X.12050.
- K. I. Jacobsen, S. E. Lester, and B. S. Halpern. A global synthesis of the economic multiplier effects of marine sectors. *Marine Policy*, 44:273–278, 2014. ISSN 0308597X. doi: 10.1016/j.marpol.2013.09.019.
- V. Kahui, C. W. Armstrong, and G. K. Vondolia. Bioeconomic analysis of habitat-fishery connections: Fishing on cold water coral reefs. *Land Economics*, 92:328–343, 2016.
- V. A. Karatayev, M. L. Baskett, D. J. Kushner, N. T. Shears, J. E. Caselle, and C. Boettiger. Grazer behaviour can regulate large-scale patterning of community states. *Ecology Letters*, 24:1917–1929, 9 2021. ISSN 14610248. doi: 10.1111/ele.13828.
- S. Kimball, M. Lulow, Q. Sorenson, K. Balazs, Y. C. Fang, S. J. Davis, M. O’Connell, and T. E. Huxman. Cost-effective ecological restoration. *Restoration Ecology*, 23:800–810, 11 2015. ISSN 1526100X. doi: 10.1111/rec.12261.
- N. Kriegisch, S. E. Reeves, E. B. Flukes, C. R. Johnson, and S. D. Ling. Drift-kelp suppresses foraging movement of overgrazing sea urchins. *Oecologia*, 190:665–677, 7 2019. ISSN 14321939. doi: 10.1007/s00442-019-04445-6.
- E. D. Lorenzo and N. Mantua. Multi-year persistence of the 2014/15 north pacific marine heatwave. *Nature Climate Change*, 6:1042–1047, 7 2016. ISSN 1758-6798. doi: 10.1038/nclimate3082.
- P. Marjoram, J. Molitor, V. Plagnol, and S. Tavaré. Markov chain monte carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 100:15324–15328, 12 2003. ISSN 00278424. doi: 10.1073/PNAS.0306899100.
- M. L. McPherson, D. J. Finger, H. F. Houskeeper, T. W. Bell, M. H. Carr, L. Rogers-Bennett, and R. M. Kudela. Large-scale shift in the structure of a kelp forest ecosystem co-occurs with an epizootic and marine heatwave. *Communications Biology* 2021 4:1, 4:1–9, 3 2021. ISSN 2399-3642. doi: 10.1038/s42003-021-01827-6.
- A. Minter and R. Retkute. Approximate Bayesian Computation for infectious disease modelling. *Epidemics*, 29, 12 2019. ISSN 18780067. doi: 10.1016/j.epidem.2019.100368.
- L. E. Morgan, L. W. Botsford, S. R. Wing, and B. D. Smith. Spatial variability in growth and mortality of the red sea urchin, *Strongylocentrotus franciscanus*, in northern California. *Canadian Journal of Fisheries and Aquatic Sciences*, 57(5):980–992, 2000. doi: 10.1139/f00-060.

- NOAA Physical Sciences Laboratory. Marine heatwaves, 2025. Image provided by the NOAA Physical Sciences Laboratory, Boulder, Colorado. Available at: <https://psl.noaa.gov/marine-heatwaves/>.
- J. Reid, L. Rogers-Bennett, F. Vasquez, M. Pace, C. A. Catton, J. V. Kashiwada, and I. K. Taniguchi. The economic value of the recreational red abalone fishery in northern California. *California Fish and Game*, 102:119–130, 2016.
- L. Rogers-Bennett and C. A. Catton. Marine heat wave and multiple stressors tip bull kelp forest to sea urchin barrens. *Scientific Reports* 2019 9:1, 9:1–9, 10 2019. ISSN 2045-2322. doi: 10.1038/s41598-019-51114-y.
- J. N. Sanchirico and M. Springborn. How to get there from here: Ecological and economic dynamics of ecosystem service provision. *Environmental and Resource Economics*, 48:243–267, 2011. ISSN 09246460. doi: 10.1007/S10640-010-9410-5.
- D. A. Smale, M. T. Burrows, A. J. Evans, N. King, M. D. Sayer, A. L. Yunnice, and P. J. Moore. Linking environmental variables with regional-scale variability in ecological structure and standing stock of carbon within UK kelp forests. *Marine Ecology-Progress Series*, 542:79–95, 1 2016. ISSN 0171-8630. doi: 10.3354/MEPS11544.
- J. G. Smith, J. Tomoleoni, M. Staedler, S. Lyon, J. Fujii, and M. T. Tinker. Behavioral responses across a mosaic of ecosystem states restructure a sea otter-urchin trophic cascade. *Proceedings of the National Academy of Sciences of the United States of America*, 118(11):e2012493118, 3 2021. ISSN 1091-6490. doi: 10.1073/pnas.2012493118.
- T. Toni, D. Welch, N. Strelkowa, A. Ipsen, and M. P. Stumpf. Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *Journal of the Royal Society Interface*, 6:187–202, 2 2009. ISSN 17425662. doi: 10.1098/rsif.2008.0172.
- B. M. Turner and T. V. Zandt. A tutorial on approximate Bayesian computation. *Journal of Mathematical Psychology*, 56:69–85, 4 2012. ISSN 0022-2496. doi: 10.1016/J.JMP.2012.02.005.
- R. Varela, L. Rodríguez-Díaz, M. de Castro, and M. Gómez-Gesteira. Influence of eastern upwelling systems on marine heatwaves occurrence. *Global and Planetary Change*, 196, 1 2021. ISSN 09218181. doi: 10.1016/j.gloplacha.2020.103379.
- G. K. Vondolia, W. Chen, C. W. Armstrong, and M. D. Norling. Bioeconomic modelling of coastal cod and kelp forest interactions: Co-benefits of habitat services, fisheries and carbon sinks. *Environmental and Resource Economics*, 75:25–48, 1 2020. ISSN 15731502. doi: 10.1007/s10640-019-00387-y.

- J. A. Vásquez, S. Zuñiga, F. Tala, N. Piaget, D. C. Rodríguez, and J. M. Vega. Economic valuation of kelp forests in northern Chile: values of goods and services of the ecosystem. *Journal of Applied Phycology*, 26: 1081–1088, 4 2014. ISSN 09218971. doi: 10.1007/s10811-013-0173-6.
- M. Ward, T. McHugh, K. Elsmore, M. Esgro, J. Ray, M. Murphy-Cannell, I. Norton, and J. Freiwald. Restoration of northern California bull kelp forests: A partnership-based approach, June 2022. URL <https://www.reefcheck.org/wp-content/uploads/2022/06/Restoration-of-Northern-California-bull-kelp-RCF-final-report-to-OPC.pdf>.
- D. Wegmann, C. Leuenberger, and L. Excoffier. Efficient approximate Bayesian computation coupled with Markov chain Monte Carlo without likelihood. *Genetics*, 182:1207–1218, 8 2009. ISSN 00166731. doi: 10.1534/genetics.109.102509.
- T. Wernberg, K. Krumhansl, K. Filbee-Dexter, and M. F. Pedersen. *Status and trends for the world’s kelp forests*, pages 57–78. Elsevier, 1 2018. ISBN 9780128050521. doi: 10.1016/B978-0-12-805052-1.00003-6.
- R. D. Wilkinson. Approximate Bayesian computation (ABC) gives exact results under the assumption of model error. *Statistical Applications in Genetics and Molecular Biology*, 12:129–141, 5 2013. ISSN 15446115. doi: 10.1515/sagmb-2013-0010.
- C. Williams, S. Rees, E. V. Sheehan, M. Ashley, and W. Davies. Rewilding the sea? A rapid, low cost model for valuing the ecosystem service benefits of kelp forest recovery based on existing valuations and benefit transfers. *Frontiers in Ecology and Evolution*, 10, 4 2022. ISSN 2296701X. doi: 10.3389/fevo.2022.642775.

Appendix

A Model derivation details

We provide detailed derivations for: (1) drift kelp consumption and urchin reproduction dynamics, and (2) kelp dynamics with exogenous spore recruitment.

A.1 Drift kelp consumption and urchin reproduction dynamics

This appendix section provides the detailed derivation of the urchin reproduction dynamics driven by drift kelp consumption, as summarized in equations (7) and (9) in the main text. In a continuous-time framework, the evolution of drift kelp, denoted by D_t , follows the law of motion:

$$\frac{dD}{dt} = o_1 \cdot X_t - o_2 \cdot D \cdot P_t - o_3 \cdot D \cdot R_t - \Delta \cdot D, \quad (1)$$

where o_1 represents the reproduction rate of drift kelp derived from adult kelp, o_2 and o_3 denote the consumption rates by purple and red urchins, respectively, and Δ embodies the loss rate of drift kelp due to ocean currents within a specific patch. Assuming that drift kelp reaches a state of equilibrium significantly faster than adult kelp, purple, and red urchins, the equilibrium condition of drift kelp (i.e., when $\frac{dD}{dt} = 0$), is given by:

$$\bar{D}_t = \frac{o_1 \cdot X_t}{o_2 \cdot P_t + o_3 \cdot R_t + \Delta}. \quad (2)$$

Furthermore, we assume that the loss rate of drift kelp due to ocean currents out of the patch is much faster than the urchin consumption (i.e., $\Delta \gg o_2 \cdot P_t + o_3 \cdot R_t$). Therefore, the drift kelp at equilibrium is:

$$\bar{D}_t = \frac{o_1 \cdot X_t}{o_2 \cdot P_t + o_3 \cdot R_t + \Delta} \approx \frac{o_1 \cdot X_t}{\Delta}. \quad (3)$$

We redefine the ratio of the drift kelp reproduction rate to the loss rate as $\frac{o_1}{\Delta} \equiv o_4$. Consequently, the approximated equilibrium state of drift kelp is expressed by $\bar{D}_t = o_4 \cdot X_t$.

For drift kelp consumption, we assume that red urchins, owing to their larger body size, are superior competitors to purple urchins and therefore graze first. The reproductive output of red urchins from consuming drift kelp, scaled by the energy conversion factor o_5 , is given by:

$$r_D(X_t, R_t) = o_5 \cdot \bar{D}_t \cdot (1 - \exp(-\tau_R \cdot R_t)) \approx o_5 \cdot o_4 \cdot X_t \cdot (1 - \exp(-\tau_R \cdot R_t)), \quad (4)$$

where τ_R denotes the rate at which the abundance of drift kelp declines per unit of grazing intensity exerted by red urchins. By redefining the product of the energy conversion factor and the ratio of drift kelp reproduction to loss rate as $o_4 \cdot o_5 \equiv \epsilon_4$, the equation for red urchin reproduction via drift kelp consumption simplifies to:

$$r_D(X_t, R_t) = \epsilon_4 \cdot X_t \cdot (1 - \exp(-\tau_R \cdot R_t)). \quad (5)$$

Subsequent to the red urchin drift kelp consumption above, purple urchins also consume drift kelp, leading to reproductive output given by: scaled by the energy conversion factor o_6

$$p_D(X_t, P_t, R_t) = \epsilon_3 \cdot X_t \cdot \exp(-\tau_R \cdot R_t) \cdot (1 - \exp(-\tau_P \cdot P_t)), \quad (6)$$

where ϵ_3 embodies the product of o_4 (the ratio of drift kelp reproduction to loss rate) and o_6 (the energy conversion factor for purple urchin reproduction from drift kelp consumption). τ_P represents the rate of impact on drift kelp abundance per unit of grazing by purple urchins.

A.2 Kelp dynamics with exogenous spore recruitment

We modify the kelp–urchin system by introducing an external spore-recruitment term η_E . The state-transition equation becomes

$$X_{t+1} = \left[(v \cdot k(X_t, P_t, R_t) + \eta_E) \cdot \frac{\exp(-\mu_P \cdot g(X_t, P_t) - \mu_R \cdot R_t)}{1 + \phi \cdot k(X_t, P_t, R_t)} + k(X_t, P_t, R_t) \cdot \delta_X \right] \cdot (1 - \xi \cdot Z_t) + \eta_X \cdot A_{X,t}.$$

Rationale and analytical derivation of exogenous spore recruitment η_E

The 2020 purple-urchin removal at Noyo, California, offers a natural experiment that isolates the external spore-recruitment parameter η_E . Because kelp (X) was virtually absent immediately after the intervention, the biomass recorded one year later must originate from exogenous spores rather than local recruitment. The derivation is summarised below.

1. **Site and intervention:** Only *purple* urchins were removed during summer–autumn 2020; red urchins remained. The monitored patch covers 360 m².
2. **Initial state ($t = 0$):** $X_0 = 0$, $P_0 = 325$, and $R_0 = 468$.
3. **One-year outcome ($t = 1$):** $X_1 = 50.5$.
4. **Model reduction:** With $X_0 = 0$, the local-recruitment term $k(X_0, P_0, R_0)$ is zero. Setting $Z_0 = 0$ and

$A_{X,0} = 0$, the transition equation collapses to

$$X_1 = \eta_E \cdot \exp(-\mu_R \cdot R_0 - \mu_P \cdot g(X_0, P_0)).$$

5. Closed-form solution:

$$\eta_E = \frac{X_1}{\exp(-\mu_R \cdot R_0 - \mu_P \cdot P_0)} = \frac{50.5}{\exp(-\mu_R \cdot 468 - \mu_P \cdot 325)} \approx 88.3,$$

where μ_R and μ_P are the grazing-rate parameters estimated in the baseline ABC calibration. Substituting these values into the closed-form expression above yields the numerical estimate $\eta_E \approx 88.3$ which represents the annual exogenous spore input required to replicate the observed kelp rebound under the measured urchin densities at Noyo.

B Stochastic dynamic programming solution method

Below, we detail the procedure used to solve the model. We adopt a grid-based approach—i.e., a discrete state space approximation to the continuous state space—and use value function iteration to identify the solution. The model is solved using *MDPsolve*, a dynamic programming package developed by Fackler (2011), implemented in the *MATLAB* environment.

1. We discretize the continuous state space for X , P , and R using uniformly spaced grid points. Since P spans a larger range, we allocate more grid points to P relative to X and R . Similarly, we discretize the control variables A_X and A_P using uniform spacing.
2. We construct a grid consisting of all possible discretized state-action space combinations,

$$M \equiv (X_t, P_t, R_t, A_x, A_p).$$
3. For computational efficiency, we remove rows in M that contain infeasible state-action combinations, ensuring the non-negativity of biomass variables. For instance, we exclude rows where P becomes negative due to urchin removal efforts (A_P).
4. We construct the Markov state transition matrix under the stochastic shock Z and its corresponding probability weight vector W . In our context, $Z = [0, 1]$ and $W = [1 - \alpha_i, \alpha_i]$ for $i \in \{l, m, h\}$. Denote the constructed Markov transition matrix as Π . This is a *column-stochastic* transition matrix, meaning that column j represents a state-action pair given the current state, and row i represents the next period state given the state-action pair in column j .
5. Denote S^+ as the next-period state given the current states (X_t, P_t, R_t) , policy actions (A_x, A_p) , and stochastic shock Z , determined through a system of state transition equations G :

$$S^+ = [X_{t+1}, P_{t+1}, R_{t+1}] = G(X_t, P_t, R_t, A_x, A_p \mid Z).$$

However, typically S^+ will not land exactly on the discretized grid. In this case, we construct linear weights for allocating probability. When S^+ falls between two adjacent grid points, we distribute probability weights accordingly. Suppose, for simplicity, that we have a single state variable. Let S_i and S_{i+1} be the two closest grid points to S^+ . The transition probability is then assigned as:

$$\chi_i = \frac{S_{i+1} - S^+}{S_{i+1} - S_i}, \quad \chi_{i+1} = \frac{S^+ - S_i}{S_{i+1} - S_i},$$

ensuring that $\chi_i + \chi_{i+1} = 1$.

6. Since the next state depends on the stochastic shock $z_q \in Z$, we compute the transition probability matrix by summing over all possible realizations of Z weighted by their probabilities $w_q \in W$:

$$\Pi_{ij} = \sum_q w_q \cdot \chi_{ij}(z_q),$$

where Π_{ij} represents the probability of transitioning from state-action pair j to state i .

7. Extrapolation may be required if S^+ falls outside the grid domain. However, extrapolation can lead to negative probabilities in the transition matrix. To correct this, we apply a normalization mechanism. Negative probabilities are truncated to zero, and the remaining probabilities are renormalized to obtain a closely related and valid probability matrix given by¹²:

$$\hat{\Pi}_{ij} = \frac{\max\left(0, \sum_q w_q \cdot \chi_{ij}(z_q)\right)}{\sum_i \max\left(0, \sum_q w_q \cdot \chi_{ij}(z_q)\right)} = \frac{\max(0, \Pi_{ij})}{\sum_i \max(0, \Pi_{ij})}.$$

8. Once the problem is solved, we can obtain: 1) the value function, 2) the policy function for A_X and A_P at each given state grid, and 3) the Markov transition matrix under the optimal policy.
9. Through the recursive multiplication of the Markov transition matrix under optimal policies with (any) initial state vector, we compute the long-run (e.g., 50 years) probability mass function for kelp and urchin levels.

¹²Normalization to construct a non-negative transition matrix might create an absorbing state at the boundaries of the state grid. However, in most well-defined problems, this normalization has little effect on finding solutions (Fackler, 2011). Furthermore, we did not observe any absorbing behavior in the dynamics in our long-run analysis.

C Model parameterization

This appendix details data sources and methods used for selecting model parameters.

C.1 Biological model parameterization

To address the complexity and parameter uncertainty inherent in modeling kelp forest dynamics, we adopt Approximate Bayesian Computation (ABC) as our primary approach for biological model parameterization. Below, we outline the rationale for selecting ABC, provide a conceptual overview of the method, and detail our specific parameterization strategy and estimation procedure.

C.1.1 Overview of Approximate Bayesian Computation (ABC)

In ecology and biology, parameter estimation has traditionally relied on likelihood-based methods. However, the structural complexity of ecological processes often makes likelihood calculations intractable, even with modern computational advances (Csilléry et al., 2010; Beaumont, 2019). To address this, Approximate Bayesian Computation (ABC) has been proposed, a simulation-based alternative. ABC defines a parameter space and specifies priors as random distributions over biological parameters and the initial states of kelp and urchin populations. Parameters are sampled from this space, applied to a dynamic model, and used to generate simulated data. This output is then compared to observational data. Parameter sets producing simulations closely matching observations are retained, yielding an approximate posterior distribution of plausible parameters.

Let ψ be a vector of parameters of our interest (i.e., $\psi = \{\gamma_P, \gamma_R, \mu_P, \mu_R, \dots, \phi, \nu\}$) and $Y_{obs} = \{y_{obs,i}\}_{i=1}^n$ be the observed data with n number of observation sites. According to Bayes' theorem, the posterior distribution of ψ given the observed data is proportional to the product of prior and likelihood:

$$Pr(\psi|Y_{obs}) \propto Pr(Y_{obs}|\psi) \cdot Pr(\psi), \quad (7)$$

where $Pr(\psi)$ is the prior and $Pr(Y_{obs}|\psi)$ is the likelihood. For a single parameter, outcome variable, and spatial location, ABC begins by specifying a prior distribution for the parameter and an initial value for the outcome. In each iteration, a parameter value is randomly drawn from the prior and used to run the dynamic model, generating a time series of simulated outcomes. These are compared to the observed outcomes, and if sufficiently close, the sampled parameter is “accepted.” Repeating this process yields a set of accepted values that approximates the posterior distribution, which can then be used for subsequent modeling.

In the ABC framework, summary statistics are used to represent the observed/simulated data in a

relatively low dimension to raw data in a high dimension in order to decide whether to accept sampled parameters (Marjoram et al., 2003; Wegmann et al., 2009; Beaumont, 2019). Let $d(s_1, s_2)$ be a distance metric function, such as the Euclidean distance measure, between two vectors (s_1, s_2) . Let $s(\cdot)$ represent a vector of summary statistics for either true or simulated data $\hat{Y}$. Further, let Λ represent a distance tolerance to accept a given sampled parameter. A sampled parameter is accepted if the summary statistics of simulated data are sufficiently close to the observed data (i.e., $d(s(Y_{obs}), s(\hat{Y})) \leq \Lambda$) otherwise, it is rejected. A large number of iterations are conducted in order to obtain a sufficient number of accepted sampled parameters. Then, we use representative values of the distribution of accepted parameters by using the median or mean of each parameter distribution. By choosing ‘sufficient’ summary statistics that represent the true posteriors of parameters ψ , the distribution of sampled parameters is ‘approximately’ close to the true posterior (Jabot et al., 2013; Arroyo-Esquivel et al., 2023; Marjoram et al., 2003; Turner and Zandt, 2012):

$$Pr(\psi|Y_{obs}) = Pr(\psi|s(Y_{obs})) \approx Pr(\psi|d(s(Y_{obs}), s(\hat{Y})) \leq \Lambda). \quad (8)$$

The conventional rejection algorithm in ABC is often inefficient, yielding low acceptance rates when the posterior differs markedly from the prior or when the parameter space is high-dimensional (Wilkinson, 2013; Toni et al., 2009; Alahmadi et al., 2020). To improve efficiency, Markov Chain Monte Carlo (MCMC) and Sequential Monte Carlo (SMC) methods have been proposed. We adopt the MCMC approach introduced by Marjoram et al. (2003). Another challenge in ABC is selecting an appropriate distance tolerance Λ for parameter acceptance. To address this, we use the automated tolerance selection algorithm proposed by Wegmann et al. (2009).

We extend the ABC approach of Arroyo-Esquivel et al. (2023) with several modifications to enhance the reliability of parameter estimation. While Arroyo-Esquivel et al. (2023) relied on data from two sites with a single interannual period (2007–2008) updated annually, we use data from six sites along the Northern California coast with six annual observations (2007–2012). For each site, we pair all consecutive observation periods (e.g., 2007–2008, 2011–2012) to maximize the sample size. Additionally, whereas Arroyo-Esquivel et al. (2023) used a single criterion—the root mean square error (RMSE) of kelp density—for one species over one period, we adopt a broader distance metric following Minter and Retkute (2019): the sum of the squared root of squared distances (SRSD), equivalent to the root squared error (RSE), applied across all three species over all observation periods:

$$D = \sum_{n=1}^6 d_n \quad (9)$$

$$d_n = \sqrt{\sum_t (X_{n,t} - \hat{X}_{n,t})^2} + \sqrt{\sum_t (R_{n,t} - \hat{R}_{n,t})^2} + \sqrt{\sum_t (P_{n,t} - \hat{P}_{n,t})^2},$$

where n indexes the sites, and $\hat{X}$, $\hat{R}$, and $\hat{P}$ are simulated densities of kelp (X), red urchins (R), and purple urchins (P) based on candidate parameters drawn from the priors. Summing SRSD values across sites into the overall distance metric D balances two goals: (1) reducing excessive parameter rejections due to heterogeneous site dynamics,¹³ and (2) capturing the dynamics of all three species effectively. Parameter vectors with $D \leq \Lambda$ are accepted as the most plausible representations of kelp forest dynamics

The estimation was executed in the R programming environment, using the *EasyABC* package. Data for this parameterization, specifically the kelp and urchin count dataset, was provided by Reef Check. The ABC estimation procedure is described below in two distinct, sequential phases.

C.1.2 Prior setup for parameters

The ABC methodology necessitates that researchers propose a plausible range of priors for each parameter within the explored space. In our model, which incorporates 14 biological parameters, 11 were directly estimated using ABC. The remaining three parameters, specifically the mortality rates of bull kelp, purple urchins, and red urchins, were derived from existing ecology and biology literature.

We established uniform priors for each of the 11 biological parameters, defined by specific minimum and maximum ranges. These ranges were determined through a combination of educated guesses, consultations with professional biologists and ecologists from the Kelp RiSES team, and preliminary experiments. These experiments were aimed at identifying the range of parameter combinations that could result in coexistence behavior within urchin and kelp dynamics. Our parameter estimation reflects the interannual dynamics of the species at the representative patch scale of this research, specifically a $360m^2$ area. This patch size corresponds to the dimensions of the count survey conducted by Reef Checks.

C.1.3 Posterior distribution estimation

Given the limited data environment, the significant uncertainty surrounding the biological processes of each species and their interactions, and the broadly uncertain prior ranges for each parameter (albeit informed by consultations with experts), our ABC approach relies on a *trial-and-error* procedure. The detailed procedure is outlined below.

1. Specify models for the biological dynamics of kelp, purple urchins, and red urchins.

¹³Using many summary statistics increases the likelihood of parameter rejections and computational complexity.

2. Pair all possible consecutive periods of species observations for each site. For instance, $(X_{1,2007}, X_{1,2008})$, where $X_{1,2007}$ is used as the initial input for site 1 in 2007 to simulate kelp dynamics and generate the output $\hat{X}_{1,2008}$. The simulated $\hat{X}_{1,2008}$ and observed $X_{1,2008}$ are then used to compute the distance measure following Equation (9). This procedure is applied identically to purple and red urchins.
3. Specify priors for each parameter based on educated guesses, literature, and professional consultations. For the initial trial, the specific priors for each parameter are as outlined. Once the priors are defined, implement ABC to sample the posterior distribution of each parameter.
4. Sample 3 million candidate parameter values in the ABC using MCMC. Retain the upper 1% of sampled values for each parameter as the posterior distribution, resulting in 30,000 retained samples.
5. Based on the sampled posterior distributions, use the median value of each parameter to solve the model and evaluate its fitness by assessing the long-run state distribution of each species under a low-MHW regime scenario, which emulates the last 50 years of the kelp forest ecosystem in California.
6. Through visual inspection and fitness evaluation measures (e.g., using MAE) across species, assess how well the sampled parameters replicate the observed biological processes of the kelp ecosystem in Northern California. For instance, this evaluation can be visualized as the figure below.




Figure C1: Comparison of probability density functions for kelp forest states based on sampled model parameters and observed data from Northern California. (Blue: computed long-run probability density based on sampled parameters; Red: observed long-run density). The computed density is obtained by applying the Markov transition matrix over a 50-year period and integrating over the other two state variables.

7. Repeat the above procedure multiple times until there is no clear improvement in parameter fitness as evaluated by the long-run distribution. In our case, a total of approximately 150 repetitions were conducted to obtain the preferred parameter sets (posterior distributions for each parameter).
8. Use the preferred sampling distributions for the rest of the analysis, such as sensitivity analysis.

C.2 Marine heatwave (MHW) probability

The annual probability of MHW occurrence for each regime (α_i , for $i = \{l, m, h\}$) was determined using a dataset of extreme water temperature days from McPherson et al. (2021). This dataset records the number of days with extreme water temperatures, as detected by satellite remote sensing along the coast of Mendocino and Sonoma County in Northern California, spanning from 1985 to 2020. Following the definition by Hobday et al. (2016), an extreme water temperature day is identified as a sequence of 5 consecutive days with temperatures above the 90th percentile based on a 30-year historical temperature distribution in the region. Analysis of the data generally reveals years characterized by either a low extreme of fewer than 40 days or a high extreme exceeding 80 days, with minimal overlap between these ranges.




Figure C2: Historic marine heatwave (MHW) days from 1985 to 2020 for Mendocino and Sonoma Counties in Northern California. Data for 1985–2019 are from McPherson et al. (2021), with 2020 supplemented using records from NOAA Physical Sciences Laboratory (2025). For 2020, the MHW duration was approximately three months, converted to days as reported for coastal regions of Mendocino County. Darker shaded bars indicate years with a high duration of MHW days.

Based on this dataset, the high MHW regime ($i = h$) probability, α_h , was calculated by counting the number of years within the period from 2010 to 2020 that experienced a spike in extreme water temperature days and dividing by the total number of years in the period. Similarly, α_l was calculated for the years 1985–2009. The medium MHW frequency parameter, α_m , is derived as the average of α_h and α_l .

C.3 Fishery, management and non-ABC derived model parameterization

We identified plausible parameters for fishery, management, and non-ABC estimation-based components using existing peer-reviewed literature, government reports, and personal communications with experts. For

the fishery and management model parameters below, we detail the information source and any transformation applied to match the scale of our representative patch. All monetary values were converted to 2023 dollars using the Consumer Price Index (CPI).

1. η_X (marginal effect of kelp outplanting intensity per hour, 2.133 kelp/hr): 32 juvenile kelp strips cultivated, processed, and transplanted within a total of 15 person-hours, as per Carney et al.(2005).
2. η_P (marginal effect of purple urchin removal per diving hour, 320.16 urchins/hr): 1) Total weight of removed purple urchins: 14,292 kg (OPC 2022). 2) Total diving hours for restoration: 720 hrs (OPC 2022). 3) Average weight of a purple urchin: 62 g/urchin (CDFW personal communication). Calculation: $14,292 \text{ kg} / 62 \text{ g} / 720 \text{ hrs} = 320.16 \text{ urchins/hr}$.
3. $\max(A_X)$, maximum kelp outplanting effort range (120 hours): Based on Duggins et al., cultured bull kelp was outplanted with spacing of 1–2 m^2 . Assuming 1 kelp stipe is outplanted for every 1.5 m^2 over a 360 m^2 patch, and at a rate of 2.1 kelps planted per hour per person, the required effort is $360m^2/1.5m^2/2.1 \text{ kelps/hr} \approx 115$ hours. Rounding up, it was set at 120 with a buffer.
4. $\max(A_P)$, maximum urchin removal effort range (10 hours): Educated guess based on the spent time per patch as reported in the kelp restoration report by Ward et al. (2022).
5. δ_X (annual natural survival rate of bull kelp, 0.15): Bull kelp is an annual species. Approximately 15% of bull kelp survives from one year to the next, based on communications with a regional partnership group.
6. δ_R (annual natural survival rate of red urchin, 0.93): Mean annual survival rate of red urchins across multiple sites in Northern California, estimated by Ebert et al. (1999).
7. δ_P (annual natural survival rate of purple urchin, 0.88): Annual survival rate of purple urchins at Point Arena in Mendocino County, estimated by Ebert (2010).
8. λ (Proportion of commercially harvestable urchins within the total population, 0.55): Age structure information taken from Ebert et al. (1999) and share calculated based on a typical commercial red urchin size suggested by CDFW (California Department of Fish and Wildlife (2020)).
9. ω (Annual fishing mortality in the red urchin fishery of Northern California, 0.4): Estimated by Ebert et al. (1999).
10. β (Social discount factor, 0.97): Assumed value.


11. ρ (Red urchin ex-vessel price, \$3.4/urchin): Based on a 3-year average price (2020–2022) California Department of Fish and Wildlife (2020), adjusted for CPI.
12. κ_R (catch revenue multiplier for adjacent industrial sectors): 3.1 per revenue dollar — This value reflects the mean multiplier effect of U.S. commercial fisheries, which typically ranges between 2 and 4. For example, in California’s tuna fisheries, the multiplier spans from 1.29 to 3.66 (Jacobsen et al., 2014). Export-oriented fisheries, such as tuna and red urchin (particularly those exported to Japan), tend to exhibit higher multiplier effects.
13. κ_X (unit direct kelp value per kelp, \$1.63/kelp): The maximum achievable red urchin fishery revenue and associated multiplier benefits—as a representative commercially valuable fishery reliant on kelp forests—is estimated at \$812 per patch. This is calculated as the ex-vessel price of red urchins multiplied by the harvestable red urchin population (i.e., fishing mortality rate times the maximum red urchin density in a given patch) along with multiplier benefits to coastal communities. Dividing this maximum fishery revenue by the maximum kelp biomass in the patch (500 kelps) yields a per-kelp value: $\$812 / 500 = \1.63 per kelp.
14. P_{exo} (Purple urchin exogenous recruitment for recolonization, 53 urchins/ $360m^2$): Based on the restoration report provided by Ward et al. (2022), purple urchin recolonization after removal is 17.6% of pre-restoration biomass. With an average pre-MHW biomass of 300 urchins/ $360m^2$ in RC data, this results in $360 \times 17.6\% \approx 53$ urchins per patch.
15. R_{exo} (Red urchin exogenous recruitment for recolonization, 25 urchins/ $360m^2$): Assuming the same recolonization rate and average biomass of 150 urchins/patch in RC data, the exogenous recruitment is 25 urchins/patch.
16. ξ (MHW shocks, 0.4): Determined through consultation with a kelp biologist (a research partnership member) through lab experiment.
17. θ_X (Kelp outplanting cost per person hour, \$42.5/hr): $\$382/15 \text{ hrs} = \$25.46/\text{hr}$ in 2002, adjusted for CPI, as per Carney et al. (2005).
18. θ_P (Purple urchin removal cost per diving hour, \$102.6): Total cost for purple urchin removal: $\$65,490 / 720 \text{ hrs} = \$90.9/\text{hr}$ adjusted to $\$102.6/\text{hr}$ post-CPI adjustment, as per Ward et al. (2022).
19. θ_H (Red urchin harvest cost per fishing trip, \$16.21/patch): (1) Based on an interview with an urchin diver in Ft. Bragg. Minimum earning per trip is \$500 with \$150 for tender operation. Ward et al. (2022) states that divers participating in purple urchin removal paid \$600 for a ship. We assume a reasonable

cost of \$600 per trip. (2) Assuming a representative patch (360 m^2) is harvested once a year and the minimum profitable harvest per trip is 500 urchins. (3) The average red urchin population of our dataset is 57 in a given patch over 2007-2012. By multiplying the fishing mortality and commercially harvestable red urchin population share estimated by Morgan et al. (2000), on average, 13 urchins are harvested from a patch. (4) By dividing 500 urchins/trip by harvested urchin per patch (i.e., 13/patch), it gives the number of patches harvested per trip, N_{patch} 4) Total red urchin harvest cost divided by the number of patches visited per trip (N_{patch}) gives the harvest cost per patch: \$16.21/patch.


D Results: supplementary figures

D.1 Value functions

D.1.1 Value function under different MHW regimes




(a) Value function (low MHW regime; exploitation stage; midpoint latent states)




(b) Value function (high MHW regime; exploitation stage; midpoint latent states)

Figure D3: Value functions for alternative scenarios (exploitation stage; midpoint latent states), shown under the low MHW regime (top row) and high MHW regime (bottom row). Each value function is shown as a heatmap for each pair of state variables on the vertical and horizontal axes.

D.1.2 Restored and degraded states




(a) Value function at the restored latent state (medium MHW regime; exploitation stage)




(b) Value function at the degraded latent state (medium MHW regime; exploitation stage)

Figure D4: Value functions for alternative scenario (medium MHW regime; exploitation stage, evaluated at the restored (top) and degraded (bottom) latent states). Each value function is shown as a heatmap for each pair of state variables on the vertical and horizontal axes.

D.1.3 Value function under restricted restoration



(a) Value function (medium MHW regime; exploitation stage; midpoint latent states; $A_P = 0$)




(b) Value function (medium MHW regime; exploitation stage; midpoint latent states; $A_X = 0$)


Figure D5: Value functions for alternative scenarios (medium MHW regime; exploitation stage; midpoint latent states), under restricted implementation of restoration actions. Top: urchin removal restricted ($A_P = 0$). Bottom: kelp outplanting restricted ($A_X = 0$). Each value function is shown as a heatmap for each pair of state variables on the vertical and horizontal axes.

D.2 Policy functions

D.2.1 Policy function under different MHW regimes




(a) Kelp outplanting intensity (low MHW regime; exploitation stage; midpoint latent states)




(b) Kelp outplanting intensity (high MHW regime; exploitation stage; midpoint latent states)

Figure D6: Optimal restoration levels for alternative scenarios (exploitation stage; midpoint latent states), shown under the low MHW regime (top row) and high MHW regime (bottom row). Each policy function is shown as a heatmap for each pair of state variables on the vertical and horizontal axes.




(a) Purple urchin removal intensity (low MHW regime; exploitation stage; midpoint latent states)




(b) Purple urchin removal intensity (high MHW regime; exploitation stage; midpoint latent states)

Figure D7: Optimal restoration levels for alternative scenarios (exploitation stage; midpoint latent states), shown under the low MHW regime (top row) and high MHW regime (bottom row). Each policy function is shown as a heatmap for each pair of state variables on the vertical and horizontal axes.

D.2.2 Policy functions evaluated at the restored and degraded states




(a) Kelp outplanting intensity (medium MHW regime; exploitation stage; degraded latent state)




(b) Kelp outplanting intensity (medium MHW regime; exploitation stage; restored latent state)

Figure D8: Optimal restoration levels for alternative scenarios (exploitation stage), shown under the medium MHW regime, evaluated at the degraded (top row) and restored (bottom row) latent states. Each policy function is shown as a heatmap for each pair of state variables on the vertical and horizontal axes.




(a) Purple urchin removal intensity (medium MHW regime; exploitation stage; degraded latent state)




(b) Purple urchin removal intensity (medium MHW regime; exploitation stage; restored latent state)

Figure D9: Optimal restoration levels for alternative scenarios (exploitation stage), shown under the medium MHW regime, evaluated at the degraded (top row) and restored (bottom row) latent states. Each policy function is shown as a heatmap for each pair of state variables on the vertical and horizontal axes.

D.2.3 Restoration stage policy function over time




(a) Kelp outplanting intensity, $A_X(X, R, P)$ at $t = 1$ (restoration stage; equal regime belief; midpoint latent states)




(b) Kelp outplanting intensity, $A_X(X, R, P)$ at $t = T$ (restoration stage; equal regime belief; midpoint latent states)

Figure D10: Optimal restoration levels for kelp outplanting intensity at the initial (top row) and final (bottom row) restoration periods (restoration stage; equal regime belief; midpoint latent states). Each policy function is shown as a heatmap for each pair of state variables on the vertical and horizontal axes.




(a) Purple urchin removal intensity, $A_p(X, R, P)$ at $t = 1$ (restoration stage; equal regime belief; midpoint latent states)




(b) Purple urchin removal intensity, $A_p(X, R, P)$ at $t = T$ (restoration stage; equal regime belief; midpoint latent states)

Figure D11: Optimal restoration levels for purple urchin removal intensity at the initial (top row) and final (bottom row) restoration periods (restoration stage; equal regime belief; midpoint latent states). Each policy function is shown as a heatmap for each pair of state variables on the vertical and horizontal axes.

D.2.4 Sided MHW belief




(a) Kelp outplanting intensity, $A_X(X, R, P)$ at $t = 15$ (restoration stage; medium MHW regime; midpoint latent states)




(b) Kelp outplanting intensity, $A_X(X, R, P)$ at $t = 15$ (high-MHW sided belief; restoration stage; medium MHW regime; midpoint latent states)

Figure D12: Kelp outplanting intensity at time $t = 15$ in the restoration stage (midpoint latent states), under two different sided beliefs about regime: low-MHW belief (top row) and high-MHW belief (bottom row). Each policy function is shown as a heatmap for each pair of state variables on the vertical and horizontal axes.

D.2.5 Restricted reward under no return of urchin harvest in the exploitation stage



(a) Kelp outplanting intensity (medium MHW regime; exploitation stage; no red urchin harvest benefit)




(b) Purple urchin removal intensity (medium MHW regime; exploitation stage; no red urchin harvest benefit)


Figure D13: Optimal restoration levels for an alternative scenario (medium MHW regime; exploitation stage; no red urchin harvest benefit), showing kelp outplanting intensity (top row) and purple urchin removal intensity (bottom row). Each policy function is shown as a heatmap for each pair of state variables on the vertical and horizontal axes.

D.3 Long-run analysis


D.3.1 Bivariate long-run probability density functions under different MHW regimes



(a) Long-run state distribution with no restoration (low MHW)




(b) Long-run state distribution with no restoration (medium MHW)




(c) Long-run state distribution with no restoration (high MHW)


Figure D14: Bivariate long-run state distributions with no restoration for the low- (top), medium- (middle), and high-MHW (bottom) regimes.



(a) Long-run state probability distributions under optimal restoration (low MHW)




(b) Long-run state distribution under optimal restoration (medium MHW)




(c) Long-run state distribution under optimal restoration (high MHW)

Figure D15: Bivariate long-run state distributions with restoration for the low- (top), medium- (middle), and high-MHW (bottom) regimes.

D.3.2 Integrated probability density functions under different MHW




(a) No restoration under low MHW




(b) No restoration under high MHW

Figure D16: Probability density functions for state variable levels after 50 years under no restoration with low MHW (top row) and high MHW (bottom row).




(a) Long-run probability density for state variable levels under low MHW with restoration




(b) Long-run probability density for state variable levels under high MHW with restoration

Figure D17: Probability density functions for state variable levels after 50 years under optimal restoration for the low-MHW (top) and high-MHW (bottom) regimes.


D.3.3 Integrated probability density functions at the end of restoration stage with inclined MHW belief



(a) Post-restoration probability density at $t = T$ (restoration stage; low-MHW sided belief; initial state: degraded)



(b) Post-restoration probability density at $t = T$ (restoration stage; equal regime belief; initial state: degraded)



(c) Post-restoration probability density at $t = T$ (restoration stage; high-MHW sided belief; initial state: degraded)

Figure D18: Probability density functions for state variable levels at $t = T$ under different regime beliefs: low-MHW (top), equal (middle), and high-MHW (bottom). In all cases, the restoration stage is assumed and the initial state is degraded.