UC Davis

UC Davis Electronic Theses and Dissertations

Title

Essays on the Economics of Coastal Communities and Marine Resources as Integrated Socio-Environmental Systems

Permalink

https://escholarship.org/uc/item/42s1m4c7

ISBN

9798193264507

Author

Kim, Kyumin

Publication Date

2026-09-08

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Essays on the Economics of Coastal Communities and Marine Resources as Integrated Socio-Environmental Systems

By

KYUMIN KIM

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Agricultural and Resource Economics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Michael R. Springborn, Co-Chair

Matthew N. Reimer, Co-Chair

James N. Sanchirico

Committee in Charge

2026

Acknowledgments

I am deeply grateful to my dissertation committee for their guidance and support throughout my doctoral studies. My co-chairs, Michael R. Springborn, who accepted me into the program and believed in my potential, and Matthew N. Reimer, who showed me seemingly endless patience, have profoundly shaped the way I think about economics, what research means, and how research is actually done. I have learned an enormous amount from both of them, including, I suspect, more patience than either knew he had. I am also grateful to James N. Sanchirico for his thoughtful advice on both research direction and career, which has been invaluable throughout my time at UC Davis. I owe a special debt of gratitude to Richard J. Sexton, who believed in me as a prospective ARE student when others were uncertain and whose unwavering support helped make this journey possible.

I have been fortunate to share this journey with wonderful friends and colleagues at UC Davis. I am especially grateful to T.J. Anderson, my dear office soulmate; Kingswell and Faith Huang, who were always wonderful companions for dinners and, laughter; Hyunjung Lee, my only Korean cohort mate; and Ed Whitney in Agricultural and Resource Economics. I am equally grateful to Mei Blundell, Ryan Swanson and Gabby Yang. The conversations, coffee breaks, seminars, dinners, and shared struggles of graduate school made these years immeasurably richer. I thank all of them for their friendship, companionship, and the many moments of joy along the way.

I am also deeply grateful to my family in Korea, especially my mother, whose love and support, shared through countless phone calls across the distance, sustained me throughout these years. Her encouragement has been a constant source of strength throughout this journey.

Finally, and most importantly, I thank my best friend and big brother from another mother, *W. Lorenzo Swank*, who has stood by me through every step of this journey, from Korea to the United States and through all the highs and lows of doctoral life. Through my hardest moments, he never wavered in his belief in me, even when I doubted myself. His friendship, encouragement, and unwavering support carried me through more moments than I can count. This dissertation would not exist without him, and it is to him that I owe my deepest gratitude.

Abstract

Coastal and fishing communities depend on marine resources whose productivity is increasingly shaped by climate-driven environmental change, exposing local economies to overlapping ecological and economic uncertainty. This dissertation adopts the framework of coupled human-environmental systems to examine how the linkage between marine resources and the communities that depend on them is sustained under such uncertainty. Spanning production-frontier estimation, dynamic bioeconomic modeling, and causal panel time-series inference, the three essays trace this linkage through three channels: (1) the sectoral structure of local economies tied to the resource sector, (2) the active management of degraded ecosystems, and (3) the transmission of harvest-sector shocks to local non-fishing economies. The first essay examines how diversification of fisheries and of the broader industrial base shapes economic growth and stability in Alaska's fishing communities. Both forms of diversification improve long-term performance, and a complementary relationship between industrial diversification and fisheries specialization allows a dual-track strategy that reduces the trade-off between growth and stability. The second essay develops a dynamic bioeconomic model that derives welfare-maximizing, state-contingent restoration strategies for a coupled kelp-urchin system under stochastic marine heatwave regimes. Optimal deployment of the two tools contrasts sharply: kelp outplanting is most valuable when kelp biomass is low, whereas urchin removal is applied aggressively across most states. Welfare and optimal policy prove far more sensitive to social and management assumptions than to biological and climate alternatives. The third essay estimates the dynamic multiplier effects of commercial fisheries on local economies, providing the first estimates of the explicit time path of local multipliers for a renewable resource. Using a panel local projections framework, it finds that fishing revenue shocks generate employment effects that build well beyond the contemporaneous period, with dynamics that differ across channels and are amplified in highly fishing-dependent communities. Together, the essays show that the resilience of coastal communities under environmental change depends jointly on the structure of local economies, the management of degraded ecosystems, and our understanding of the long-run linkage between the resource sector and local economies.

Contents

Acknowledgments **ii**

Abstract **iii**

List of Figures **xiii**

List of Tables **xiv**

1 Introduction **1**

2 The Role of Fisheries Diversification in Economic Growth and Stability: Evidence from Alaska's Fishing Economies **5**

Abstract 5

2.1 Introduction 6

2.2 Empirical setting and data 8

2.3 Diversification, growth, and growth instability in Alaskan communities 10

2.4 A hyperbolic distance function approach 15

2.5 Results 21

2.5.1 Robustness check 1: Testing endogeneity of input variables 26

2.5.2 Robustness check 2: Estimation with direct diversification measures 29

2.5.3 Robustness check 3: Estimation with alternative diversification measures using the Herfindahl–Hirschman index 29

2.6 Discussion 30

Appendices **33**

A Estimation procedure details in stochastic frontier analysis 34

B Measuring economic growth rate and economic instability 35

C Robustness check using direct diversification measure 37

D Robustness check using alternative diversification measure with inverse HHI index 38

E Constructing 95% confidence intervals for MP_{y,x_3,x_4} 39

3 Balancing Multiple Interventions for Dynamically Efficient Kelp Forest Restoration Under Marine Heatwave Uncertainty 42

Abstract 42

3.1 Introduction 43

3.2 Bioeconomic model of population dynamics and optimal management in the kelp-urchin system 46

3.2.1 Biological dynamics 46

3.2.2 Management components 52

3.2.3 The management problem 53

3.2.4 Parameterization 54

3.3 Results 54

3.3.1 Baseline scenario 56

3.3.2 Biology and climate scenarios 60

3.3.3 Social and management scenarios 64

3.4 Discussion 69

Appendices 72

A Model derivation details 73

A.1 Drift kelp consumption and urchin reproduction dynamics 73

A.2 Kelp dynamics with exogenous spore recruitment 74

A.3 Fixed proportional policy: model modification and evaluation 76

B Stochastic dynamic programming solution method 78

C Model parameterization 80

C.1 Biological model parameterization 80

C.2 Marine heatwave (MHW) probability 85

C.3 Fishery, management and non-ABC derived model parameterization 86

D Results: supplementary figures 89

D.1 Alternative scenario 1: Latent states 89

D.2 Alternative scenario 2: MHW-frequency levels and uncertainty . . . 91
D.3 Alternative scenario 3: Kelp spore loss and exogenous recruitment . . . 98
D.4 Alternative scenario 4: Static proportional heuristic policy . . . 99
D.5 Alternative scenario 5: No return of the red urchin harvest . . . 100
D.6 Alternative scenario 6: Restricted restoration . . . 101

4 Dynamic Multiplier Effects of Commercial Fisheries in Local Fishing Economies 103

Abstract . . . 103
4.1 Introduction . . . 104
4.2 Methods and Data . . . 107
4.2.1 Local Projections with Instrumental Variables (LP-IV) . . . 107
4.2.2 Approximate Bias Correction for Dynamic Panel Local Projections . . . 110
4.2.3 Monte Carlo Simulation . . . 112
4.2.4 Data . . . 113
4.3 Results . . . 115
4.3.1 Spillover Effects: Resident Earnings versus Local Landings . . . 115
4.3.2 Sectoral Decomposition: Traded and Non-traded Employment . . . 119
4.3.3 Heterogeneity by Fishing Dependency . . . 121
4.3.4 Joint Inference on the Impulse Response Path . . . 123
4.3.5 Scaling Elasticities to Level Effects . . . 126
4.4 Discussion . . . 126

Appendices 131

A Incidental Parameter Problem (IPP) and Instrumental Variable (IV) Approach . . . 132
B Nickell Bias in Dynamic Panels under Finite T and Large N . . . 133
C Detailed Derivation of Approximate Bias Correction . . . 134
D Bias Reduction by Controlling for Past Information in Local Projections . . . 135
D.1 Setup . . . 136
D.2 Adding One Lag of y . . . 136
E Monte Carlo Simulation Design and Results . . . 138
E.1 Structural Data-Generating Process . . . 138

E.2 Time-Series SVAR(1) . . . 139
E.3 Panel SVAR(1) . . . 140
E.4 Panel SVAR(2) . . . 141
E.5 Simulation Results . . . 143
F System GMM Estimation for Joint Inference . . . 147
G Summary Statistics . . . 150
H Scaled Impulse Response Functions . . . 150

5 Conclusion **153**

Bibliography **169**

List of Figures

2.1 Employment Growth, Instability, Industrial and Fisheries Diversification in Alaskan Fishing Communities (2000-2016) 13

2.2 Economic growth–instability portfolio profiles of Alaskan fishing communities by diversification quantiles (2000–2016). Colors and shapes of dots represent industrial and fisheries diversification quantiles (1st–4th). 14

2.3 Conceptualized hyperbolic distance function characterizes the production possibility frontier between economic outputs and specialization (inverse diversification). Red dots (like Community A) represent communities that lie on the frontier—indicating fully efficient economic growth (appropriate growth and instability) at a given specialization level—while black dots (like Community B) indicate communities with room for further improvement with either more growth or less instability at given specialization. 16

2.4 Marginal products of industrial and fisheries specialization with complementary effects. Dashed lines denote 95% confidence intervals computed via the delta method (Appendix E). Estimates are based on specification (3) in Table 2.2 and are evaluated at the sample means of the other covariates. We use the HDF specification for interpretability because the EHDF case normalizes input x. 31

3.1 Value function for the baseline scenario specified in Table 3.2, shown as a heatmap for each pair of state variables on the vertical and horizontal axes, with latent states at their midpoint. 57

3.2 Optimal restoration levels (labor hours) for the two management options, kelp outplanting (top row) and purple urchin removal (bottom row), for the baseline scenario, as specified in Table 3.2. 58

3.3 Probability mass 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, conditional on the given policy, over 50 periods and integrating over the other two state variables. The starting state was our midpoint state case ($\mathcal{X} = 50, \mathcal{P} = 50, \mathcal{R} = 50$). 60

3.4 Probability density functions for state variable levels after 50 years under the baseline scenario (gray) and under the addition of kelp spore loss and exogenous recruitment (blue), given optimal restoration policy. 64

3.5 Optimal restoration proportions for the two management options: kelp outplanting (top row) and purple urchin removal (bottom row). Values are converted from restoration effort hours into proportions relative to the corresponding population levels in the baseline scenario (except latent states in plotting are degraded). Regions with $X = 0$ or $P = 0$ are excluded from the plot (shown in white), because the ratio is undefined at zero population. 65

3.6 Optimal kelp outplanting level given baseline model except with no return of red urchin harvest. 67

3.7 Percent change in the value function under restricted restoration relative to the baseline scenario: outplanting only (top, $A_P = 0$) and purple urchin removal only (bottom, $A_X = 0$). . 68

3.8 Optimal restoration levels given the baseline model except with restricted restoration: outplanting only (top, $A_P = 0$) and removal only (bottom, $A_X = 0$). 69

A1 Share of surviving adult kelp (green), grazing share of adult kelp by purple urchins (purple), and by red urchins (red), across kelp population levels under a Type IV functional response. Each subfigure corresponds to a different fixed urchin population level: (i) the mean population density observed in the long-run analysis without restoration, (ii) the midpoint population level, and (iii) the degraded state. 75

A2 Relative contributions of direct adult kelp grazing and drift kelp grazing by red urchins across kelp population levels. Each line corresponds to a different fixed purple urchin population level. 75

C3 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. 84

C4 Estimated probability density functions for kelp forest states based on sampled model parameters, literature, and observed data from Northern California with external kelp spore recruitment ($\eta_E = 23$, and $\zeta = 0.05$) (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. 85

C5 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. 86

D6 Value functions given the baseline model except with restored (top) and degraded (bottom) latent states. 89

D7 Optimal kelp outplanting intensity (labor hours) given the baseline model except with degraded (top) and restored (bottom) latent states. 90

D8 Optimal purple urchin removal intensity (labor hours) given the baseline model except with degraded (top) and restored (bottom) latent states. 90

D9 Value functions given the baseline model except with low MHW regime (top) and high MHW regime (bottom). 91

D10 Optimal kelp outplanting intensity (labor hours) given the baseline model except with low MHW regime (top) and high MHW regime (bottom). 92

D11 Optimal kelp outplanting intensity (labor hours) given the baseline model except with unknown MHW regime with low-frequency-skewed belief (top) and high-frequency-skewed belief (bottom). 93

D12 Optimal purple urchin removal intensity (labor hours) given the baseline model except with low MHW regime (top) and high MHW regime (bottom). 94
D13 Bivariate long-run state distributions with no restoration given the baseline model except with low- (top), medium- (middle), and high-MHW (bottom) regimes. 95
D14 Bivariate long-run state distributions with optimal restoration given the baseline model except with low- (top), medium- (middle), and high-MHW (bottom) regimes. 96
D15 Probability density functions for state variable levels after 50 years under no restoration given the baseline model except with low MHW (top) and high MHW (bottom). 97
D16 Probability density functions for state variable levels after 50 years under optimal restoration given the baseline model except with low-MHW (top) and high-MHW (bottom) regimes. . . 97
D17 Value function given the baseline model except with external recruitment. 98
D18 Optimal kelp outplanting (top) and purple urchin removal (bottom) intensity (labor hours) given the baseline model except with external recruitment. 98
D19 Difference in X_{t+1}, $\Delta X_{t+1} \equiv X_{t+1}^{\text{exo}} - X_{t+1}^{\text{no}}$, under external spore recruitment versus no external recruitment, evaluated under no restoration ($A_X = A_P = 0$). The black contour indicates $\Delta X_{t+1} = 0$. 99
D20 The value function showing the expected present value of long-run net benefits (color bar) under the baseline model but where restoration levels are determined by the static proportional policy. The starting state is given by the degraded scenario ($\mathcal{X} = 25$, $\mathcal{P} = 75$, $\mathcal{R} = 25$). The asterisk (*) indicates the levels of b_X and b_P that result in the maximum value across the static proportional policy alternatives. Outside the grid considered in Figure D20, values decline monotonically with increasing (b_X, b_P), eventually becoming negative (management costs outweigh any returns). 99
D21 Value function given the baseline model except with no red urchin harvest. 100
D22 Optimal kelp outplanting (top) and purple urchin removal (bottom) intensity (labor hours) given the baseline model except with no red urchin harvest. 100
D23 Value functions given the baseline model except with restricted restoration ($A_X = 0$, top) and additionally with external recruitment (bottom). 101
D24 Optimal restoration given the baseline model except with restricted tool choice and external kelp recruitment: outplanting only (top, $A_P = 0$) and removal only (bottom, $A_X = 0$). 102

4.1 LP-OLS versus LP-IV impulse response of wages and employment to resident earnings. Left panels show the baseline specification; right panels show the no-control specification. Blue lines denote LP-OLS; green lines denote LP-IV. Shaded areas represent 90% confidence intervals. 116

4.2 LP-IV versus bias-corrected (LP-IV-BCC) impulse response of wages and employment to resident earnings (baseline specification, first-stage F-statistic = 90.12). Green solid line denotes LP-IV; red dashed line denotes LP-IV-BCC following Herbst and Johannsen (2024). Shaded areas represent 90% confidence intervals. 117

4.3 LP-IV versus bias-corrected (LP-IV-BCC) impulse response of wages and employment to local landings revenue (baseline specification, first-stage F-statistic = 37.82). Green solid line denotes LP-IV; red dashed line denotes LP-IV-BCC following Herbst and Johannsen (2024). Shaded areas represent 90% confidence intervals. 118

4.4 LP-IV versus bias-corrected (LP-IV-BCC) impulse response of traded and non-traded sector employment to resident earnings (baseline specification, first-stage F = 115.1 and 98.81, respectively). Green solid line denotes LP-IV; red dashed line denotes LP-IV-BCC following Herbst and Johannsen (2024). Shaded areas represent 90% confidence intervals. 120

4.5 LP-IV versus bias-corrected (LP-IV-BCC) impulse response of traded and non-traded sector employment to local landings revenue (baseline specification, first-stage F = 31.15 and 37.39, respectively). Green solid line denotes LP-IV; red dashed line denotes LP-IV-BCC following Herbst and Johannsen (2024). Shaded areas represent 90% confidence intervals. 121

4.6 LP-IV-BCC impulse response of wages and employment to resident earnings by fishing dependency. Subsamples progressively restrict to communities above the indicated percentile threshold of the dependency distribution (community-level mean ratio of fishing revenue to total wages). Dots represent point estimates; vertical bars represent 90% confidence intervals. 122

E.1 Impulse response functions (IRFs) estimated under different specifications: levels vs. long-difference outcomes, with and without baseline controls ($p = q = 2$ lags of outcome and treatment). 143

E.2 Panel impulse response functions (IRFs) across different specifications: levels vs. long differences, with and without baseline controls ($p = q = 2$). 144

E.3 Panel impulse response functions (IRFs) under an AR(2) specification: levels vs. long differences, with and without baseline controls ($p = q = 2$). 145
E.4 Bias-corrected impulse response functions (IRFs) with baseline controls ($p = q = 2$), following Herbst and Johannsen (2021, 2024). 146
F.1 Comparison of individual LP-IV and system GMM point estimates for resident-based outcomes. Green dots denote LP-IV estimates using horizon-specific samples; blue dots denote system GMM estimates using the common sample. Vertical bars represent 90% confidence intervals. 148
F.2 Comparison of individual LP-IV and system GMM point estimates for delivery-based outcomes. Green dots denote LP-IV estimates using horizon-specific samples; blue dots denote system GMM estimates using the common sample. Vertical bars represent 90% confidence intervals. 149
H.1 Scaled cumulative impulse response functions, resident specification (LP-IV-BCC). 151
H.2 Scaled cumulative impulse response functions, delivery specification (LP-IV-BCC). 152

List of Tables

2.1 Summary statistics 12

2.2 Estimation results: (enhanced) hyperbolic distance function 22

2.3 Test results for the normalized $MRT_{s,y}$ 23

2.4 Test results for the normalized MP_{y,x_3} 23

2.5 Test results for the normalized MP_{y,x_4} 24

2.6 Test results for normalized MP_{s,x_3} 24

2.7 Test results for normalized MP_{s,x_4} 25

2.8 Results from Bayesian Mundlak–Chamberlain device estimation. 28

C1 Estimation results for robustness check: using fisheries diversification 37

D1 Estimation results for robustness check: using an alternative diversification measure (inverse HHI) 38

3.1 Parameter definitions, values, and sources 55

3.2 Model specification alternatives 56

3.3 Summary of results across alternative scenarios relative to the baseline. 61

4.1 Joint Wald test for the full impulse response path ($h = 0, \ldots, 5$) under the common-sample GMM specification (effective period: 2006–2016). 124

G.1 Summary statistics for the estimation sample at $h = 0$ (baseline specification). 150

Chapter 1

Introduction

Coastal communities have long stood at the intersection of complex ecological and economic systems. Their livelihoods depend on marine resources whose abundance, distribution, and productivity are shaped by environmental processes that operate at scales far larger than any individual community or fishery. In recent decades, the speed and severity of environmental change in coastal and marine systems have accelerated. Marine heatwaves have become more frequent and intense, reshaping species distributions and triggering ecological regime shifts. Fishery stocks have shifted across management boundaries, and entire fisheries have collapsed or been temporarily closed in response to ecosystem disruption. These changes do not act on ecosystems alone—they propagate into the communities, industries, and households that depend on marine resources for their economic well-being.

A central challenge for resource economics is to understand how coastal communities and the marine systems they depend on co-evolve under environmental and economic uncertainty, and what economic strategies, policy instruments, and management interventions can sustain community welfare and ecosystem function over the long run. Meeting this challenge requires analytical frameworks that treat humans and ecosystems not as separate spheres, but as components of a single coupled system in which feedbacks between ecological dynamics and economic decisions determine system trajectories.

This dissertation adopts the framework of *integrated socio-environmental systems (ISES)* to study three sets of questions about how coastal communities interact with marine resources under uncertainty. In an ISES framing, ecological dynamics and human decisions co-determine the state of the system: environmental shocks reshape the resource base available to economic actors, while human decisions about extraction, restoration, and economic structure feed back into ecological outcomes. Although the three essays in this dissertation address distinct research questions and employ different analytical methods, they share this common framing and a common concern with two cross-cutting issues:

1. **Addressing environmental and ecological uncertainty in fisheries and the coastal communities that rely on them.** Coastal economies face overlapping sources of uncertainty—climate-driven eco-

logical shifts, volatile fishery returns, and broader market fluctuations—and each essay examines a distinct strategy for managing or characterizing this uncertainty. Chapter 3 develops a dynamic bioeconomic model in which state-contingent restoration of kelp forests serves as an active management response to stochastic marine heatwave regimes, sustaining both ecological function and economic value under climate risk. Chapter 2 examines how a more diversified pattern of resource use—across fisheries and across broader industrial sectors—shapes community economic growth and stability, allowing fishing-dependent communities to absorb environmental and market variability while sustaining long-run growth. Chapter 4 turns to the question of how fishery shocks, once they occur, propagate through local economies over time, providing empirical evidence on the magnitude and persistence of risk transmission to employment and wages in fishing-dependent communities.

2. **Coupling between ecological and economic outcomes.** Each essay treats ecological state variables and economic outcomes as jointly determined rather than studying either in isolation. Chapter 3 explicitly models the coupling between kelp–urchin population dynamics and the economic benefits and costs of restoration interventions, deriving restoration policies that internalize ecological feedbacks and translate ecological recovery into measurable economic value. Chapter 4 estimates the long-run multiplier effects through which fishery sector activity is coupled to non-fishery employment and wages in local economies, tracing how shocks originating in the fishery propagate into the broader community over multiple years. Chapter 2 links the structure of resource use—fishery portfolio composition together with broader industrial diversification—to community-level employment growth and stability, characterizing how diversified versus specialized patterns of resource use shape community development outcomes.

Together, the three essays show that coastal community resilience and welfare under environmental change depend on the interaction of *ex-ante* economic structure, *active* ecological management, and *ex-post* economic adjustment. No single channel is sufficient; a fuller account of resilience requires considering all three.

Overview of the chapters

Chapter 2: The role of fisheries diversification in economic growth and stability. The first essay examines how the structure of local economies—specifically, the diversification of fisheries and of the broader industrial base—shapes long-run economic growth and stability in Alaska's fishing communities. Using 17 years of data covering 177 communities, I adapt a hyperbolic distance function (HDF) from production economics to characterize the joint relationship between average employment growth and its stability across communities with different portfolio structures. The empirical analysis identifies a complementary relationship between industrial diversification and fisheries specialization: communities can achieve favorable combinations of growth and stability through a "dual-track" strategy that balances diversification in the broader economy against specialization in the fisheries sector. These findings show that the *ex-ante* structure of local economic portfolios is a key determinant of community resilience to environmental and economic variability, and suggest that one-size-fits-all diversification prescriptions overlook important interactions between sectors.

Chapter 3: Balancing multiple interventions for dynamically efficient kelp forest restoration under marine heatwave uncertainty. The second essay develops a dynamic bioeconomic model that derives welfare-maximizing, state-contingent restoration strategies for a coupled kelp–urchin system under stochastic marine heatwave regimes. The model integrates kelp–urchin feedback dynamics, two restoration tools (kelp outplanting and purple urchin removal), restoration costs and benefits, and stochastic environmental forcing. Solving the stochastic dynamic programming problem yields contrasting optimal deployment patterns: kelp outplanting is most valuable when kelp biomass is low and is scaled back as the system recovers, whereas purple urchin removal is aggressively applied across most system states. The state-contingent optimal policy yields substantially higher net benefits than commonly used heuristic strategies, and restoration remains economically justified even under pessimistic assumptions about the recovery of the commercial red urchin fishery. By treating restoration as an investment under uncertainty in a coupled biological–economic system, this chapter shows how *active ecological management* can be designed to support both ecosystem function and human livelihoods over the long run.

Chapter 4: Dynamic multiplier effects of commercial fisheries in local fishing economies. The third essay quantifies the long-run multiplier effects of commercial fisheries on non-fishery employment and wages in Alaska's fishing communities. Despite the central role that commercial fisheries play in many coastal economies, evidence on how fishery sector performance propagates through local economies over multiple years remains limited. Using residence-based employment and wage data for approximately 150 fishing communities over 2000–2016, I employ a panel local projections (LP) framework with instrumental variables to trace impulse response functions of non-fishery outcomes to exogenous changes in fishery revenue. To address the incidental parameter problem that arises in dynamic panels with finite time horizons and persistent outcomes, the analysis applies recent bias-correction methods and validates their performance through Monte Carlo simulation calibrated to the data structure. The estimates show economically meaningful long-run spillovers, with effects that vary by sector and by community fishing dependence.

Contributions and roadmap

Together, the three essays advance our understanding of how local economic structure, ecological restoration, and economic propagation jointly determine the welfare of coastal communities under environmental change, and they produce findings of direct relevance to fisheries management agencies, restoration practitioners, and coastal policy makers. The remainder of the dissertation is organized as follows. Chapter 2 studies the role of fisheries and industrial diversification in shaping growth and stability in Alaska's fishing communities. Chapter 3 develops a dynamic bioeconomic model of kelp forest restoration under marine heatwave uncertainty. Chapter 4 estimates dynamic multiplier effects of commercial fisheries on local fishing economies. Chapter 5 synthesizes findings across the three essays and discusses their implications for research and policy.

Chapter 2

The Role of Fisheries Diversification in Economic Growth and Stability: Evidence from Alaska's Fishing Economies

Abstract

Climate change poses serious challenges to the growth and stability of fishing communities. While a diverse portfolio of fisheries has been shown to stabilize income for individual fishers and communities, evidence regarding its role in broader local economies remains limited. This study addresses this gap by examining how fisheries and industrial diversification shape economic growth and stability. Using a hyperbolic distance function (HDF), we analyze 17 years of fisheries and economic data for 177 Alaskan fishing communities. Our findings show that both fisheries and industrial diversification significantly affect employment growth and the stability of local fishing economies. Economic stability often comes at the cost of reduced growth, reflecting the risk-return trade-offs inherent in diversification. We also identify a complementary relationship between industrial diversification and fisheries specialization. This dual-track strategy, which balances diversification and specialization across fisheries and industrial sectors, enables communities to adapt their economic structures to local circumstances while minimizing trade-offs between growth and stability. These results suggest that policymakers should prioritize strategies that foster balanced diversification and specialization, tailored to community conditions, to strengthen resilience in the face of climate-related uncertainties.

Keywords: Fisheries diversification; Industrial diversification; Economic growth and instability; Climate resilience; Hyperbolic distance function; Stochastic frontier analysis; Alaska.

2.1 Introduction

Climate-induced changes in the ocean increasingly threaten the sustainability of coastal communities that rely on fisheries. Shifts in the distribution and productivity of fish stocks compound existing pressures from overexploitation, sea level rise, ocean acidification, and the globalization of the seafood trade (Anderson et al., 2018; Pinsky et al., 2018; Free et al., 2019; Rogers et al., 2019). These challenges are particularly acute in vulnerable communities, where declining and aging populations exacerbate risks of economic and social instability (Himes-Cornell and Kasperski, 2015; Colburn et al., 2016; Blasiak et al., 2017). As traditional fishing practices face growing uncertainties, the resilience of these communities depends on their ability to adapt to changing coastal, demographic, and economic conditions (Birkenbach et al., 2023).

Diversifying fishing activities across multiple fisheries is a promising strategy for coastal communities to adapt to the uncertainties of a changing marine environment. Just as a diversified portfolio of financial assets reduces overall return variability (Markowitz, 1952), a well-diversified portfolio of fisheries revenue streams can lower a community's exposure to climate-induced shocks, particularly when shocks affect fish populations and markets asynchronously. Extensive research supports this stabilizing effect, showing that diversification can buffer fluctuations in commercial fisheries and the livelihoods they support (Kasperski and Holland, 2013; Sethi et al., 2014; Holland et al., 2017; Young et al., 2019; Fisher et al., 2021; Abbott et al., 2023). Consequently, many studies emphasize harvesting diverse fisheries and adapting portfolios to changing conditions as a critical means of stabilizing local economies that depend on fisheries (Cline et al., 2017; Bris et al., 2018; Salgueiro-Otero et al., 2022).

While diversifying fisheries portfolios can support adaptation, it entails notable challenges. Most fisheries are not open-access and require fishers to purchase permits, which can be prohibitively expensive (Holland and Kasperski, 2016). Limited access to capital further constrains fishers' ability to invest in permits and expand fishing opportunities (Olson, 2011). Diversification may also dilute the knowledge and efficiencies gained through specialization, potentially reducing overall returns (Anderson et al., 2017). Geographic and ecological constraints matter as well: in regions with fewer or more volatile local fisheries, communities have limited options for constructing a diversified portfolio, leaving them less able to capture diversification benefits (Sethi et al., 2014; Young et al., 2019). In addition, as fisheries management has become more prescriptive to address overfishing and overcapacity, the scope for diversifying across fisheries has narrowed, and diversification has declined in many parts of the world (Holland et al., 2017; Abbott et al., 2023).

An important uncertainty in assessing the buffering role of fisheries diversification concerns the position of commercial fishing within the broader local economy. In many regions, the economic prominence of commercial fishing is declining as coastal communities face increasing competition for limited offshore space from aquaculture, tourism, offshore wind energy, and resource extraction (FAO, 2024). Future resilience will likely hinge on communities' ability to adapt to a shifting economic landscape shaped by emerging marine industries. Just as fisheries diversification can stabilize fishing revenues, industrial diversification can provide broader economic stability by offsetting losses in one sector with gains in others (Chandra, 2002, 2003; Kluge, 2018; Hafner, 2019). Diversifying across sectors reduces reliance on fisheries and mitigates exposure to fluctuations in fish populations and seafood markets. Communities that successfully diversify their economies and adjust industrial portfolios toward emerging marine sectors are therefore more likely to achieve greater economic stability and growth (van Putten et al., 2016).

Empirical evidence on the roles of fisheries and industrial diversification in the stability and growth of local fishing economies remains limited. Prior research has largely examined the stabilizing effect of fisheries diversification on fishing revenues, with less attention to local economic outcomes such as employment or wages (Sethi et al., 2014; Cline et al., 2017; Young et al., 2019; Fisher et al., 2021). Moreover, many studies emphasize reduced instability while overlooking implications for returns or growth, despite the well-established trade-off between stability and growth in modern portfolio theory (Markowitz, 1952). Examining the joint dynamics of economic growth, instability, and diversification is therefore essential for a holistic understanding of these trade-offs in coastal economies and for informing policies that support adaptation to a changing marine environment.

In this study, we examine how fisheries and industrial diversification contribute to long-term growth and stability in local fishing economies. We focus on Alaska's fishing communities, which accounted for 60% of U.S. domestic seafood landings in 2020 (NOAA, 2022). These communities face significant climate-related challenges, including warming ocean temperatures, ocean acidification, shifting fish-stock distributions, and increased price volatility (Himes-Cornell and Hoelting, 2015; Kasperski and Holland, 2013). Using 17 years of employment and fishing revenue data from 177 communities, we estimate the trade-off between employment growth and stability with a hyperbolic distance function (HDF) adapted from production economics. The HDF identifies the frontier of efficient combinations of employment growth and stability and, in turn, allows us to characterize how fisheries and industrial diversification shape efficient growth–stability outcomes. This framework provides new insight into the growth–stability relationship and

the interaction between fisheries and industrial diversification in fishing communities.

Our findings show that both fisheries and industrial diversification significantly shape employment growth and stability. Economic stability often comes at the cost of lower growth, reflecting the risk–return trade-offs inherent in diversification. We also identify a complementary relationship between industrial diversification and fisheries specialization. A dual-track strategy that balances specialization and diversification across fisheries and industrial sectors enables communities to adapt their economic structures to local circumstances while minimizing growth–stability trade-offs. Our results suggest that policymakers should prioritize strategies that foster balanced diversification and specialization across fisheries and industrial sectors to build resilience in fishing communities facing climate-related uncertainties.

2.2 Empirical setting and data

Alaska offers a valuable context for studying how fisheries and industrial diversification shape local economic growth and stability. The state's commercial fishing sector is substantial, generating $5.2 billion in wholesale revenues in 2022 and employing 48,000 workers, representing 66% of manufacturing jobs (ASMI, 2024). Its economic impact is significant: a 10% increase in fishery earnings raises resident income by 0.7% (i.e., $1.54 per additional dollar earned) (Watson et al., 2021). At the same time, Alaska's economy is not solely dependent on fishing. Since the 1980s, the oil and gas sector has been dominant, and tourism, shipping, transportation, and the public sector also play vital roles. These industries often surpass fishing as primary sources of employment in coastal communities, underscoring the importance of industrial diversification for local economies.

Alaska is also a compelling case study because unique data measure employment and wages by workers' place of residence. Traditional datasets, such as the Bureau of Labor Statistics' Quarterly Census of Employment and Wages (QCEW), typically report employment and wage statistics by place of work and at aggregated geographies such as the county. This approach presents two challenges. First, county-level statistics can obscure meaningful empirical relationships, particularly in regions like Alaska and the Western U.S., where counties are geographically large and encompass heterogeneous communities. Second, discrepancies between where individuals live and work complicate inferences about people rather than places (Guettabi and James, 2020; Jacobsen et al., 2023).

The Alaska Local and Regional Information (ALARI) database addresses these challenges by reporting

employment and wage data by industrial sector measured at workers' places of residence. This approach is enabled by the Alaska Department of Labor and Workforce Development's (AKDOL) linkage of unemployment insurance records (also used in the QCEW) with Alaska's Permanent Fund Dividend (PFD) application data. The PFD, paid annually to eligible residents who have lived in the state for a full calendar year and intend to remain indefinitely, includes applicants' residential addresses. This linkage provides granular, community-level insight into employment and wage patterns and offers a basis for examining how fisheries and industrial diversification relate to local economic growth and stability.

We compile a comprehensive dataset of economic and commercial fishing variables for all Alaskan communities involved in commercial fisheries between 2000 and 2016. Data on permit-owner earnings are sourced from the Alaska Commercial Fisheries Entry Commission (CFEC) Basic Information Tables, which provide annual records of harvests and earnings for each community–fishery pair. Alaskan fisheries are defined by species, fishing district, and gear type, and participation requires a fishery-specific permit issued by the CFEC. In 2010, 20,275 permits were issued across 205 fisheries in Alaska, with permit-owner communities identified by the addresses listed on their permits. The value of deliveries to local processors is aggregated from individual delivery records reported through the Alaska Department of Fish and Game's (ADF&G) fish tickets and eLandings systems.

Information on local wages and employment comes from the ALARI database, which provides sector-specific economic data for 344 Alaskan communities over the same period. Complementing these data, demographic variables, such as average population size, were obtained from the Alaska Bureau of Labor Statistics to give broader context to the local economies of fishing communities. For our analysis, we focus on 177 communities that consistently engaged in commercial fishing throughout the study period (2000-2016), covering a period of 17 years.

It is important to clarify that ALARI records are derived from unemployment insurance data, which exclude many categories of workers in the seafood industry. Commercial fishers and crew, who are typically self-employed or contract workers, are not captured. Wages for upstream and downstream proprietors and other self-employed individuals are also omitted. By contrast, wage and employment data for workers employed by commercial processors are included in ALARI. This distinction is crucial for interpreting which types of employment are affected by fisheries and industrial diversification. Accordingly, we do not consider the direct effect of fisheries diversification on fisheries employment.

A further consideration is that ALARI measures the number of individual workers rather than full-time

equivalent (FTE) jobs. Because Alaska's seafood industry is highly seasonal, many workers are employed only for a few months each year. Our employment estimates may therefore differ from studies that report FTE measures.

2.3 Diversification, growth, and growth instability in Alaskan communities

In this study, we use employment growth rather than GDP to measure economic growth because employment data are available at a finer resolution (both geographically and by industry), whereas GDP is typically reported only at broader geographies (e.g., state). Employment metrics therefore provide more detailed micro-level insights into local fishing economies and align with our focus on industrial diversification, as sector-level job trends capture economic shifts more consistently. We measure a community's average economic growth using the geometric growth return of employment, which captures compounding (as with compound interest) and provides a more accurate view of long-run outcomes. Following Lande (1994), Chandra (2002), and Hafner (2019), this measure is obtained by estimating: [1]

$$\log(Emp_{t,c}) = \beta_{0,c} + \beta_{1,c} \cdot t + \epsilon_{t,c}, \tag{2.1}$$

where $Emp_{t,c}$ is employment in community c at year t, and $\beta_{1,c}$ denotes the instantaneous employment growth rate for community c. We convert this to an average annual geometric growth return, $1 + \mu_c = \exp\{\beta_{1,c}\}$, where μ_c is community c's geometric growth rate. Economic instability is quantified as the standard deviation of the residuals from the ordinary least squares (OLS) estimate of the log-linear time trend in Equation (2.1). The fitted trend captures the long-run (geometric) return. Larger positive or negative deviations from this trend are interpreted as instability in year-to-year growth. Further details and justification are provided in Appendix B.

Previous studies commonly measure diversification using the inverse Herfindahl–Hirschman Index (HHI) or the Shannon diversity index (Kasperski and Holland, 2013; Holland and Kasperski, 2016; Sethi et al., 2014; Kluge, 2018; Hafner, 2019). We adopt the Shannon diversity index, defined as the exponential of Shannon entropy, to quantify industrial diversification. This measure has intuitive interpretation: if a community achieves the same level of economic growth across N industrial sectors and then doubles the number to $2N$ while maintaining the same growth, the index also doubles. It represents the effective number of perfectly

[1] This approach is also employed in reports by the World Bank (Chandra, 2002).

balanced sectors within a community (Jost, 2006; Abbott et al., 2023).

We calculate this index from employment *growth* across 13 industrial sectors within a community, rather than from employment levels.[2] Formally,

$$Div_c = \frac{1}{T}\sum_{t=1}^{T}\exp\left(-\sum_{i=1}^{N} P_{ict}\,\ln P_{ict}\right), \tag{2.2}$$

$$P_{ict} = \frac{|\Delta Emp_{ict}|}{\sum_{j=1}^{N}|\Delta Emp_{jct}|}, \tag{2.3}$$

where P_{ict} is the share of total employment growth contributed by sector i in community c in year t, and ΔEmp_{ict} is sector-i employment growth. The Shannon-based diversification measure in Equation (2.2) equals the exponential of Shannon entropy evaluated on $\{P_{ict}\}_{i=1}^{N}$. It approaches 1 when a single sector dominates and increases with diversification; Div_c averages this annual index over the T years.

Kluge (2018) advocates growth-based diversification measures over traditional size-based measures (i.e., shares). Size-based indices can misleadingly encourage regions dominant in stable sectors to diversify into more volatile ones, potentially reducing stability rather than enhancing it; they also fail to distinguish how positive versus negative growth affects stability. As a result, size-based measures may mask underlying instability during periods of rapid sector-specific expansion or contraction, even when the industrial composition appears unchanged. A growth-based approach more accurately captures variability in sector contributions and provides a more reliable indicator of economic stability.

For fisheries diversification, we follow Sethi et al. (2014) and also compute a Shannon diversity index based on fishing permits. Although many studies use realized fishing revenues or landings, those metrics are distorted by stochastic catch variation and price volatility. Using the number of active permits provides a more stable, ex-ante measure of diversification that avoids noise from short-run fluctuations in catch volumes and market conditions.

Table 2.1 reports summary statistics for the variables used in this study, calculated as averages over the study period. Long-run economic growth, measured by the average annual geometric return of employment in local fishing economies, is close to one, indicating near-zero growth. On average, employment in Alaska's fishing communities was largely stable or slightly declining, with an annual decrease of approximately 0.3%.

Industrial diversification averages 5.7 effective, perfectly balanced sectors based on the North American

[2]Our local fishing economy data classify industrial sectors using North American Industry Classification System (NAICS) codes, which are also used for federal and state labor statistics.

Variable	N	Mean	SD	Min	Max
Employment Growth (Return)	177	0.997	0.019	0.919	1.060
Growth Instability	177	1.080	0.062	1.008	1.387
Industrial Diversification	177	5.747	1.881	1.354	8.996
Fisheries Diversification	177	3.601	3.596	1.000	19.962
Population (Individuals)	177	2,908	21,457	13	282,781
Wage per Capita (USD)	177	9,998	4,313	1,325	27,362
Fishing Permit Share within Borough	177	0.151	0.246	0.002	1.000
Fishing Earnings to Wage Income Ratio	177	0.397	0.955	0.000	7.998

Notes: All values are averages over the study period.

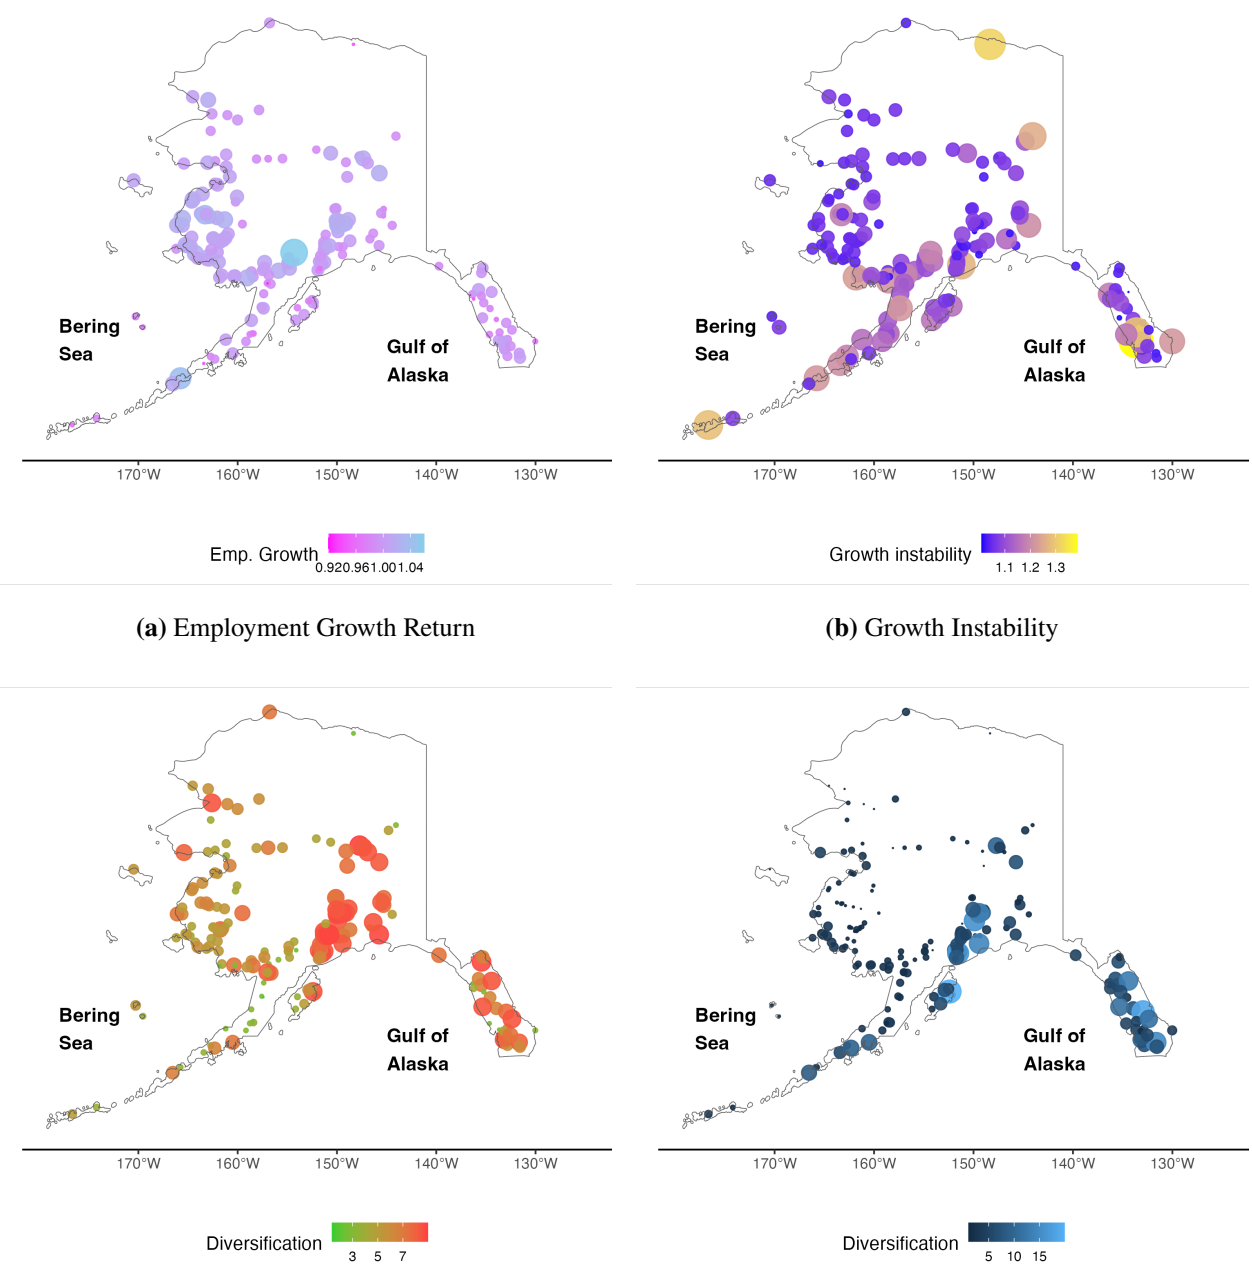
Table 2.1: Summary statistics

Industry Classification System (NAICS). Fisheries diversification indicates typical engagement in 3–4 distinct fisheries per community, with notable heterogeneity across communities.

Population size across Alaskan fishing communities is highly dispersed, ranging from small settlements such as Ugashik (roughly a dozen residents on average) to Anchorage, the largest population center. This dispersion suggests that most fishing communities are moderately to sparsely populated. Fishing permit shares (a measure of spatial fisheries concentration within a borough) indicate that most communities are not dominant fishing hubs, whereas a few—including Sitka, Yakutat, Anchorage, and Juneau—emerge as hotspots with concentrated permits. The fishing earnings-to-wage income ratio, a measure of fishing dependency, points to moderate reliance on fishing: on average, revenue from the sector accounts for about 40% of total wage income from other industries, although dependence varies substantially, as reflected by the high standard deviation.

Figure 2.1 presents spatial patterns and relationships among the main variables. Alaska's fishing communities are generally concentrated along the coast in the Southwest, Southcentral (around the Gulf of Alaska), and Southeast regions, with some inland communities along the Yukon River primarily for salmon fisheries. Over the study period, employment growth was relatively steady overall, with higher growth in the Southwest than elsewhere. In Southcentral, communities with higher instability are primarily along the Gulf of Alaska, and these areas overlap with higher economic growth; however, in general, high growth does not strongly coincide with high instability.

Industrial and fisheries diversification are positively correlated across communities. Communities in



(a) Employment Growth Return

(b) Growth Instability

(c) Industrial Diversification

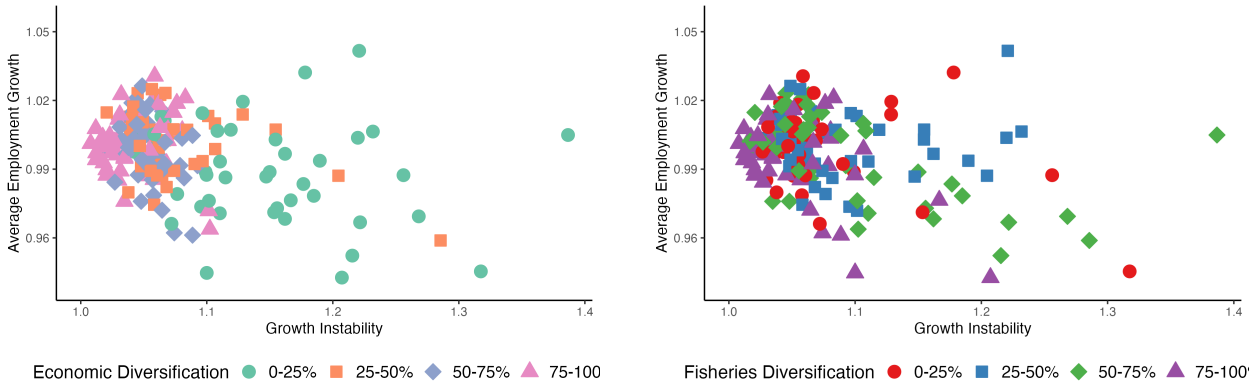
(d) Fisheries Diversification

Figure 2.1: Employment Growth, Instability, Industrial and Fisheries Diversification in Alaskan Fishing Communities (2000-2016)

Western Alaska display low diversification in both industrial structure and fisheries yet are associated with higher economic growth. This pattern suggests that more specialized regions (e.g., the Southwest) may be driving growth. Finally, the relationship between diversification (both industrial and fisheries) and economic instability is less distinct but generally negative. Communities along the Aleutian Islands (Bering Sea) and

in Southeast with less diversification exhibit higher instability, consistent with the high-risk, high-return implication of modern portfolio theory.

Figure 2.2 provides a closer look at the relationship between employment growth return, economic growth instability, and diversification (industrial or fisheries). By plotting economic instability on the x-axis, employment growth on the y-axis, and coloring each fishing community by diversification quantile (in 4 groups, from low to high), we further examine their relationship in detail.



(a) By industrial diversification quantiles

(b) By fisheries diversification quantiles

Figure 2.2: Economic growth–instability portfolio profiles of Alaskan fishing communities by diversification quantiles (2000–2016). Colors and shapes of dots represent industrial and fisheries diversification quantiles (1st–4th).

In the left panel of Figure 2.2, the growth-instability portfolio for fishing communities is shown by industrial diversification. The figure provides a couple of insights: First, the distribution of the fishing communities' risk-return portfolio appears to form a general minimum-variance frontier as suggested by MPT. This indicates that the relationship between economic growth, as measured by employment in Alaska fishing communities, and the associated economic instability is inherently non-linear.

Previous research has analyzed risk and return in fishing revenue, typically focusing on the risk-diversification relationship using linear risk measures like the Coefficient of Variation (CV) of fishing revenue (Sethi et al., 2014; Kasperski and Holland, 2013; Abbott et al., 2023; Gokhale et al., 2024). However, if a non-linear relationship between economic growth and instability exists, a linear approach may fail to capture the true dynamics of risk-return-diversification. To address this limitation, our estimation approach (described below) relaxes the linearity assumption and employs a more flexible functional form to better understand the nuanced relationship between a community's economic growth and instability.

Second, the figure illustrates that fishing communities with higher industrial diversification (i.e., 3rd (50-

75%) and 4th quantiles (75-100%)) tend to cluster near zero growth (i.e., close to one as return) but with lower instability. This indicates that communities with greater industrial diversification may, on average, experience a steady-state employment growth status. Conversely, communities with low industrial diversification, particularly those in the first and second industrial diversity quantiles, exhibit greater dispersion in the growth-instability space.

2.4 A hyperbolic distance function approach

We adopt the hyperbolic distance production function (HDF) approach, commonly used in production economics. In this framework, economic growth is treated as a desirable output, while economic risk is considered an undesirable output in the growth process of communities. Here, x represents a vector of K inputs ($x \in \mathbb{R}_+^K$), y denotes a vector of V desirable outputs ($y \in \mathbb{R}_+^V$), and s refers to a vector of U undesirable outputs ($s \in \mathbb{R}_+^U$), produced as byproducts of y. The production technology is defined by the production possibility set:

$$T = \{(x, y, s) : x \in \mathbb{R}_+^K, y \in \mathbb{R}_+^V, s \in \mathbb{R}_+^U, x \text{ can produce } (y, s)\}. \tag{2.4}$$

The production technology is also characterized by a distance function, which provides insights into production dynamics by analyzing three key relationships: (1) the shape of the production frontier, including the curvature between outputs ($y - y'$ for desirable outputs and $y - s$ for desirable and undesirable outputs); (2) the relationship between inputs and outputs ($y - x$), analogous to a production function; and (3) the complementarity or substitutability between inputs ($x - x'$), such as the marginal rate of technological substitution (MRTS) along the isoquant (Morrison-Paul et al., 2000; Cuesta et al., 2009; Dalheimer et al., 2024). This approach offers a comprehensive framework for studying growth and stability dynamics in fishing communities, considering diversification and other community-specific factors.

The traditional output distance function measures the maximum feasible expansion of only desirable outputs (y) required to reach the boundary of the production possibility set (T). It evaluates the efficiency of a production bundle relative to the production frontier, defined as:

$$D^O(x, y, s) = \inf_\theta \{\theta > 0 : (x, \frac{y}{\theta}, s) \in T\}. \tag{2.5}$$

When $D^O = 1$, the production bundle is efficient and lies on the frontier. Conversely, $D^O < 1$ indicates

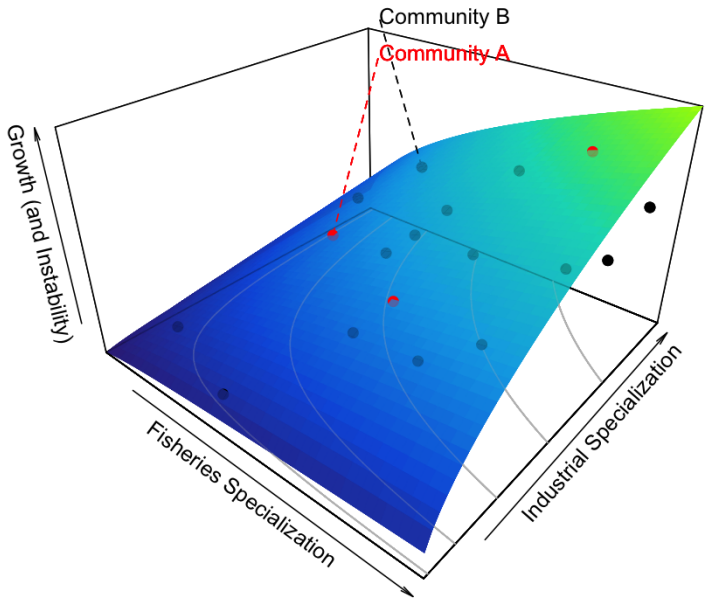


Figure 2.3: Conceptualized hyperbolic distance function characterizes the production possibility frontier between economic outputs and specialization (inverse diversification). Red dots (like Community A) represent communities that lie on the frontier—indicating fully efficient economic growth (appropriate growth and instability) at a given specialization level—while black dots (like Community B) indicate communities with room for further improvement with either more growth or less instability at given specialization.

inefficiency, suggesting that expanding desirable outputs (y) can bring the production bundle closer to the frontier. However, this function does not differentiate between desirable and undesirable outputs, limiting its capacity to address trade-offs between these two types of outputs effectively.

To overcome this limitation, the hyperbolic distance function (HDF) is introduced. The HDF allows for simultaneous adjustment of desirable and undesirable outputs, enabling a more nuanced measurement of production efficiency. By capturing both economic growth and instability as desirable and undesirable outputs, respectively, in our context, the HDF provides a comprehensive framework for evaluating the economic growth process in fishing communities. The HDF is defined as:

$$D^H(x, y, s) = \inf_\theta\{\theta > 0 : (x, \frac{y}{\theta}, \theta s) \in T\}, \tag{2.6}$$

where the HDF takes values between 0 and 1 (i.e., $D^H \in [0, 1]$). This function satisfies the following properties: (1) non-decreasing in desirable outputs, $D^H(x, \kappa y, s) \leq D^H(x, y, s)$ for $\kappa \in [0, 1]$; (2) non-increasing in undesirable outputs, $D^H(x, y, \kappa s) \leq D^H(x, y, s)$ for $\kappa \geq 1$; (3) non-increasing in inputs, $D^H(\kappa x, y, s) \leq$

$D^H(x, y, s)$ for $\kappa \geq 1$; and (4) almost homogeneity conditions, $D^H(x, \kappa y, \kappa^{-1} s) = \kappa D^H(x, y, s)$ (Cuesta et al., 2009). Using a flexible translog function, the HDF for the economic production process of community c can be specified as:

$$\begin{aligned}\ln D_c^H(x, y, s) = \alpha_0 &+ \sum_{k=1}^{K} \alpha_k \ln x_{k,c} + \frac{1}{2} \sum_{k=1}^{K} \sum_{l=1}^{K} \alpha_{kl} \ln x_{k,c} \ln x_{l,c} + \sum_{v=1}^{V} \beta_v \ln y_{v,c} + \frac{1}{2} \sum_{v=1}^{V} \sum_{l=1}^{V} \beta_{vl} \ln y_{v,c} \ln y_{l,c} \\ &+ \sum_{u=1}^{U} \delta_u \ln s_{u,c} + \frac{1}{2} \sum_{u=1}^{U} \sum_{l=1}^{U} \delta_{ul} \ln s_{u,c} \ln s_{l,c} + \sum_{k=1}^{K} \sum_{v=1}^{V} \gamma_{kv} \ln x_{k,c} \ln y_{v,c} + \sum_{k=1}^{K} \sum_{u=1}^{U} \eta_{ku} \ln x_{k,c} \ln s_{u,c} \\ &+ \sum_{v=1}^{V} \sum_{u=1}^{U} \theta_{vu} \ln y_{v,c} \ln s_{u,c},\end{aligned} \tag{2.7}$$

Suppose a function F satisfies almost homogeneity of degree (k_1, k_2, k_3, k_4), then F satisfies the following condition:

$$k_1 \sum_{k=1}^{K} \frac{\partial F}{\partial x_k} + k_2 \sum_{v=1}^{V} \frac{\partial F}{\partial y_v} + k_3 \sum_{u=1}^{U} \frac{\partial F}{\partial s_u} = k_4 F \tag{2.8}$$

Since the hyperbolic distance function must satisfy almost homogeneity of degree $(0, 1, -1, 1)$, and using the partial derivatives derived in Equation (2.7), the following condition should be satisfied for the HDF:

$$\sum_{v=1}^{V} \frac{\partial \ln D^H}{\partial \ln y_v} - \sum_{u=1}^{U} \frac{\partial \ln D^H}{\partial \ln s_u} = 1. \tag{2.9}$$

To impose the homogeneity condition, we normalize the HDF by imposing the almost homogeneity condition, $D^H(x, \kappa y, \kappa^{-1} s) = \kappa D^H(x, y, s)$. Suppose $\kappa = \frac{1}{y_1}$, with y_1 being the first desirable output in the output vector y. The hyperbolic distance function then satisfies the following equation:

$$D^H(x, \frac{y}{y_1}, s \cdot y_1) = \frac{D^H(x, y, s)}{y_1}, \tag{2.10}$$

By taking the natural log on both sides,

$$\ln D_c^H(x, y^*, s^*) = \ln D_c^H(x, y, s) - \ln y_1, \tag{2.11}$$

where $y^* = y/y_1$ and $s^* = s \cdot y_1$. These values are plugged into Equation (2.7) to replace y and s. Thus, let $\ln D_c^H(x, y^*, s^*) = TL(x, y^*, s^*)$, where $TL(\cdot)$ is the translog hyperbolic distance function defined in

Equation (2.7). Furthermore, $\ln D_c^H$ is a one-sided community inefficiency component for each community c in economic production, rearranged as $\ln D_c^H = -u_c$ (Brümmer et al., 2002; Cuesta and Zofío, 2005; Cuesta et al., 2009). By rearranging the above equation with additional noise from estimation, we derive the estimation equation as follows:

$$-\ln y_1 = TL(x, y^*, s^*) + v_c + u_c, \tag{2.12}$$

where v_c is estimation noise that is normally distributed around zero, capturing all unobserved factors beyond the control of each fishing community in the economic production process. u_c represents the distance between the observed output vector and the boundary of the production possibility set. We assume u_c to be half-normally distributed around zero, as is common in the literature with this approach (Cuesta and Zofío, 2005; Cuesta et al., 2009; Zhang and Ye, 2015). Additionally, we allow u_c to be heteroskedastic as a function of community-specific characteristics (but not direct growth-relevant inputs), following Dalheimer et al. (2024) and Peña et al. (2018). To avoid the bias associated with the traditional two-step estimation process, where the frontier and inefficiency equations are estimated separately, we opt for the simultaneous estimation of the efficient frontier and inefficiency equations using maximum likelihood estimation (MLE), as outlined by Belotti et al. (2013) (Kluge, 2018; Wang and Schmidt, 2002; Belotti et al., 2013).

Post-estimation, following the procedure of Reimer et al. (2017) and Cuesta et al. (2009), with the estimated parameters and technical efficiency, we are able to compute the marginal rate of transformation (MRT) between undesirable and desirable outputs from the production, which represents the opportunity cost of reducing bad output concerning the forgone good output. The MRT is computed as a ratio of the first derivatives of the HDF with respect to desirable output y (here, employment growth) and undesirable output s (here, growth instability) as follows:

$$MRT_{s,y} = \frac{\partial s}{\partial y} = -\frac{\partial D^H / \partial y}{\partial D^H / \partial s} = -\frac{\varepsilon_{D,y}}{\varepsilon_{D,s}} \cdot \frac{s}{y}, \tag{2.13}$$

where $\varepsilon_{D,y}$ and $\varepsilon_{D,s}$ are distance function elasticities for y and s, estimated by differentiating the translog distance function.

The MRT can be challenging to interpret directly, as it depends on the specific output ratio. To clarify the trade-off in relative terms, independent of absolute output scales, Morrison-Paul et al. (2000) and Cuesta et al. (2009) suggest normalizing the MRT by the output ratio, s/y. Following this approach, we compute the

normalized MRT (NMRT), which provides insights into the relative economic growth-risk relationship in Alaskan fishing communities, independent of output scale. The NMRT is expressed as $\text{NMRT}_{s,y} = -\frac{\varepsilon_{D,y}}{\varepsilon_{D,s}}$, where the distance elasticity for the desirable output y, $\varepsilon_{D,y}$, is recovered using the homogeneity condition in the HDF context, $\varepsilon_{D,y} = 1 + \varepsilon_{D,s}$.

Similarly, we derive the marginal product (MP) of an input x_k for the desirable output y using the distance function:

$$MP_{y,x_k} = \frac{\partial y}{\partial x_k} = -\frac{\partial D^H / \partial x_k}{\partial D^H / \partial y} = -\frac{\varepsilon_{D,x_k}}{\varepsilon_{D,y}} \cdot \frac{y}{x_k}. \tag{2.14}$$

where ε_{D,x_k} is a distance function elasticity of an input x_k. As with the NMRT, the marginal product is computed with normalization by the output-input ratio and evaluated at means. This normalized marginal product (NMP) provides a straightforward interpretation of the relative contribution of input to output, independent of absolute scale. The NMP of input x on output y is expressed as $NMP_{y,x_k} = -\frac{\varepsilon_{D,x_k}}{\varepsilon_{D,y}}$.

We consider one desirable output, average annual employment growth return ($y_c = \exp(\beta_{1,c})$) and one undesirable output, economic instability (s_c), measured by the standard deviation of annual geometric employment growth over the study period for each community c, derived from the regression equation (2.1).

The input vector includes: (1) population (x_1), representing labor input or human capital for community growth (Himes-Cornell and Hoelting, 2015); (2) wage income per capita (x_2), serving as a proxy for economic development or productivity, reflecting economic capital quality and welfare (Kluge, 2018); (3) an industrial diversification measure (x_3); and (4) a fisheries diversification measure (x_4).[3]

To meet the HDF framework's requirements and enable intuitive interpretation, however, we invert these diversification measures into specialization measures ($1/x_3$ and $1/x_4$). Unlike traditional inputs such as labor and capital, diversification aims for stable, moderate growth with lower risk, aligning with portfolio theory. This unconventional input-output relationship, characterized by lower returns and risks, can complicate interpretation in the HDF context. Inverting diversification into specialization allows us to analyze how specialization influences economic growth and stability while ensuring compliance with the HDF's input properties.

Fisheries-related variables, not directly tied to broader local economy growth—such as (1) the ratio of fishing revenue to wage income (a measure of fishing dependency) and (2) fishing permit share within

[3]We limit the inclusion of variables to a maximum of four input variables that strongly relate to economic growth in fishing communities, based on available community-level data in the HDF. This limitation arises from (1) the exponential increase in the number of parameters with additional input variables in the translog specification with limited degrees of freedom in the data, and (2) potential difficulties in likelihood function convergence as the number of estimated parameters increases.

a shared economic zone (a fishing hotspot (or concentration) index at the borough level)—are included in inefficiency terms.[4] These variables enable heteroskedastic inefficiency, influencing distance function values and promoting efficient economic growth in fishing communities. By reducing unnecessary instability and fostering growth, they capture community-specific characteristics driving heterogeneous inefficiency.

To impose the homogeneity condition, we use the desirable output y_c as the normalizing factor, following Cuesta and Zofío (2005); Cuesta et al. (2009) and Reimer et al. (2017). The resulting normalized distance function is expressed as:

$$\begin{aligned} -\ln y_c = \Big(\alpha_0 + \sum_{k=1}^{4} \alpha_k \ln x_{k,c} + \frac{1}{2}\sum_{k=1}^{4}\sum_{l=1}^{4} \alpha_{kl} \ln x_{k,c} \ln x_{l,c} + \delta_1 \ln s_c^* \\ + \frac{1}{2}\delta_{11} \ln s_c^* \ln s_c^* + \sum_{k=1}^{4} \eta_{k1} \ln x_{k,c} \ln s_c^*\Big) + v_c + u_c, \end{aligned} \tag{2.15}$$

where $s_c^* = s_c \cdot y_c$. The hyperbolic distance function (HDF) in Equation (2.15) allows for efficient output adjustments. Extending this, we additionally use the Enhanced Hyperbolic Distance Function (EHDF) used by Cuesta and Zofío (2005); Cuesta et al. (2009); Dalheimer et al. (2024), which additionally adjusts inputs for greater efficiency. The EHDF, denoted as D^E, normalizes inputs by the desirable output, yielding:

$$\begin{aligned} -\ln y_c = \Big(\alpha_0 + \sum_{k=1}^{4} \alpha_k \ln x_{k,c}^* + \frac{1}{2}\sum_{k=1}^{4}\sum_{l=1}^{4} \alpha_{kl} \ln x_{k,c}^* \ln x_{l,c}^* + \delta_1 \ln s_c^* \\ + \frac{1}{2}\delta_{11} \ln s_c^* \ln s_c^* + \sum_{k=1}^{4} \eta_{k1} \ln x_{k,c}^* \ln s_c^*\Big) + v_c + u_c, \end{aligned} \tag{2.16}$$

where $x_{k,c}^* = x_{k,c} \cdot y_c$.

Normalizing inputs and outputs by the homogeneity condition reduces potential simultaneity bias, as outputs appear on both the left and right sides of the equation. This approach allows the use of output ratios as exogenous values since they represent radial expansion while holding input levels constant (Coelli and Perelman, 1996; Coelli, 2000; Cuesta and Orea, 2002). In the EHDF framework, desirable outputs are affected directly by error terms, while inputs (and undesirable outputs) are inversely affected, permitting exogenous treatment of desirable-undesirable output ratio and input-output products normalized by desirable outputs (Cuesta and Zofío, 2005; Dalheimer et al., 2024). For convenient elasticity evaluation at the means

[4] We assume that these variables may directly influence the fishing sector's productivity but not directly impact economic growth and instability, except through their effects on the fishing sector. Therefore, we do not include them in the frontier equation.

and improved convergence in maximum likelihood estimation (MLE), we normalize all variables by their geometric mean, as recommended by Cuesta et al. (2009) and Reimer et al. (2017).

2.5 Results

We estimate the hyperbolic distance function (HDF) and the enhanced hyperbolic distance function (EHDF) in Equations (2.15) and (2.16) using the four inputs defined above. We implement stochastic frontier analysis (SFA) under three specifications that vary the inefficiency determinants: (1) no controls, (2) spatial fishing permit share within the shared economic zone, and (3) the fishing revenue-to-wage income ratio to assess how fishing dependence affects community efficiency.

Table 2.2 reports estimates for HDF (1)–(3) and EHDF (1^*)–(3^*). Across all specifications, the inputs are statistically significant in first- and/or second-order terms, supporting their inclusion. The signs and significance of the coefficients on average population (α_1), wage income per capita (α_2), industrial specialization (α_3), and fisheries specialization (α_4) are consistent with the HDF's non-increasing property in inputs; in practice, coefficients are significantly negative or not statistically different from zero. Per capita wage income, a proxy for community-specific economic development and productivity, also affects economic growth. Although its first-order effect is generally not significant, higher-order terms are significant in most specifications, suggesting increasing returns to scale in development and productivity.

Economic instability, treated as an undesirable output, also conforms to the non-increasing property: estimates for δ_1 are negative and statistically significant. This pattern is consistent with the high-risk, high-return implication of modern portfolio theory (MPT). We further corroborate this trade-off by testing the output-ratio normalized marginal rate of transformation, $NMRT_{s,y}$, evaluated at sample means (Equation (2.13)); results appear in Table 2.3.

	(1)	(2)	(3)	(1*)	(2*)	(3*)
	HDF			**EHDF**		
Population (α_1)	-0.018***	-0.018***	-0.019***	-0.018***	-0.018***	-0.019***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Wage/Capita Inc. (α_2)	0.002	0.001	-0.000	0.001	0.000	-0.001
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Ind. Specialization (α_3)	-0.014*	-0.015*	-0.018**	-0.012	-0.013*	-0.016**
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Fish. Specialization (α_4)	-0.007***	-0.007***	-0.007***	-0.006***	-0.006***	-0.006***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Population sq. (α_{11})	0.008***	0.008***	0.009***	0.008***	0.008***	0.008***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Wage/Capita Inc. sq. (α_{22})	-0.021	-0.023*	-0.027**	-0.020	-0.022	-0.026*
	(0.046)	(0.014)	(0.014)	(0.014)	(0.014)	(0.013)
Ind. Specialization sq. (α_{33})	-0.001	0.001	-0.005	-0.018	-0.016	-0.021
	(0.045)	(0.045)	(0.045)	(0.044)	(0.044)	(0.044)
Fish. Specialization sq. (α_{44})	0.009**	0.009*	0.010**	0.010**	0.009**	0.010**
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)
Population * Wage/Capita Inc. (α_{12})	-0.007	-0.006	-0.007	-0.005	-0.004	-0.005
	(0.005)	(0.004)	(0.005)	(0.004)	(0.004)	(0.004)
Population * Ind. Special. (α_{13})	0.005	0.005	0.001	0.002	0.002	-0.002
	(0.008)	(0.009)	(0.009)	(0.008)	(0.008)	(0.009)
Population * Fish. Special. (α_{14})	0.008***	0.008***	0.008***	0.007***	0.007***	0.008***
	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
Wage/Capita Inc. * Ind. Special. (α_{23})	-0.016	-0.020	-0.027	-0.024	-0.025	-0.030
	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)
Wage/Capita Inc. * Fish. Special. (α_{24})	-0.003	-0.003	-0.004	-0.001	-0.001	-0.002
	(0.054)	(0.055)	(0.054)	(0.054)	(0.054)	(0.053)
Ind. Special. * Fish. Special. (α_{34})	0.015	0.015	0.014	0.016*	0.016*	0.014
	(0.010)	(0.010)	(0.009)	(0.010)	(0.010)	(0.010)
Econ. Instability (δ_1)	-0.302***	-0.296***	-0.298***	-0.321***	-0.314***	-0.317***
	(0.053)	(0.041)	(0.039)	(0.040)	(0.039)	(0.039)
Econ. Instability sq. (δ_{11})	0.281	0.435	0.369	0.617	0.743	0.664
	(0.912)	(0.929)	(0.942)	(0.860)	(0.875)	(0.886)
Population * Econ. Instability (η_{11})	-0.011	-0.016	-0.016	0.007	0.011	0.007
	(0.042)	(0.041)	(0.041)	(0.040)	(0.040)	(0.040)
Wage/Capita Inc. * Econ. Instability (η_{21})	-0.416***	-0.397***	-0.425***	-0.349***	-0.335***	-0.364***
	(0.115)	(0.116)	(0.114)	(0.114)	(0.113)	(0.112)
Ind. Special. * Econ. Instability (η_{31})	-0.073	-0.091	-0.128	-0.014	-0.022	-0.059
	(0.167)	(0.167)	(0.167)	(0.163)	(0.164)	(0.163)
Fish. Special * Econ. Instability (η_{41})	-0.063	-0.054	-0.061	-0.036	-0.025	-0.028
	(0.050)	(0.051)	(0.051)	(0.050)	(0.050)	(0.050)
Inefficiency						
Broader Permit Share		-0.720	-1.005		-0.717	-1.024*
		(0.734)	(0.645)		(0.703)	(0.618)
Fishing Earnings to Wage Income Ratio			0.051**			0.052**
			(0.025)			(0.023)
Observations	177	177	177	177	177	177

Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 2.2: Estimation results: (enhanced) hyperbolic distance function

Our main variables of interest—industrial specialization ($1/x_3$) and fisheries specialization ($1/x_4$) (as inverse diversification)—demonstrate their relevance to economic growth. Notably, fisheries specialization shows its potential role in influencing the economic growth of local economies in fishing communities beyond the fishing sector. To further clarify how these specialization measures affect community economic growth and growth instability, we tested the marginal product using Equation (2.14), based on the estimated distance function elasticities of specialization measures, economic instability, and economic growth via the implicit function theorem.

	(1)	(2)	(3)	(1*)	(2*)	(3*)
$NMRT_{s,y}$	2.315***	2.381***	2.350***	2.009***	2.064***	2.018***
	(0.442)	(0.457)	(0.442)	(0.379)	(0.392)	(0.376)

Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 2.3: Test results for the normalized $MRT_{s,y}$

For computing the marginal product on employment growth, the distance function elasticity of economic growth ($\epsilon_{D,y}$) was recovered using the homogeneity condition in Equation (2.10). Table 2.4 below presents the test results of the computed marginal product of industrial specialization on the employment growth of fishing communities. The overall positive and statistically significant marginal product confirms this relationship between specialization and growth. From the perspective of industrial diversification, this finding suggests that diversification may reduce economic returns. Overall, the statistical significance of the marginal product of industrial specialization on economic growth demonstrates that specialization drives higher economic growth in fishing communities.

	(1)	(2)	(3)	(1*)	(2*)	(3*)
NMP_{y,x_3}	0.020*	0.021*	0.025**	0.020	0.021	0.026*
	(0.011)	(0.011)	(0.012)	(0.013)	(0.013)	(0.013)

Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 2.4: Test results for the normalized MP_{y,x_3}

As shown in Table 2.5, results for fisheries specialization are similar and, in many cases, exhibit higher statistical significance. The computed marginal product of fisheries specialization, MP_{y,x_4}, is strongly significant across all specifications, with a magnitude generally 2–2.5 times lower than that of industrial specialization. Although its effect is smaller than that of industrial specialization, the same interpretation applies: fisheries specialization contributes to higher economic growth in fishing communities beyond the fishing sector. Overall, these findings indicate that specialization—whether in the broader economy or the

fishing sector—yields higher economic returns in the local economies of Alaskan fishing communities. Equivalently, diversification (the inverse of specialization) in either fisheries or the local economy reduces economic growth.

	(1)	(2)	(3)	(1*)	(2*)	(3*)
NMP_{y,x_4}	0.010***	0.010***	0.010***	0.009***	0.009***	0.010***
	(0.003)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)

Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 2.5: Test results for the normalized MP_{y,x_4}

We also test the marginal products of industrial and fisheries specialization with respect to the undesirable output, economic instability. These marginal products, $MP_{s,x_3} = \frac{\partial s}{\partial x_3}$ and $MP_{s,x_4} = \frac{\partial s}{\partial x_4}$, are obtained via the chain rule as the product of the output trade-off and the effect of specialization on growth:

$$MP_{s,x_k} = \frac{\partial s}{\partial y} \cdot \frac{\partial y}{\partial x_k} = MRT_{s,y} \cdot MP_{y,x_k}, \quad k \in \{3, 4\},$$

because a change in an input first affects the desirable output y (through MP_{y,x_k}), which in turn affects the undesirable output s (through $MRT_{s,y}$).

	(1)	(2)	(3)	(1*)	(2*)	(3*)
NMP_{s,x_3}	0.049*	0.051*	0.059**	0.038	0.043*	0.052**
	(0.027)	(0.028)	(0.028)	(0.026)	(0.026)	(0.026)

Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 2.6: Test results for normalized MP_{s,x_3}

The test results for MP_{s,x_3}—the marginal product of industrial specialization on economic instability—generally show a statistically significant positive relationship. Thus, while more industrially specialized fishing communities in Alaska benefit from higher economic growth, they must also contend with greater instability. Interpreting specialization inversely as diversification, communities with more diversified local economies experience greater stability, albeit with a moderate reduction in growth. We also test the marginal product of fisheries specialization; results in Table 2.7 are consistent with those for industrial specialization.

In sum, both types of specialization (and, inversely, diversification)—in fisheries and in the broader industrial structure—affect local economic performance: specialization raises growth but is associated with higher instability, whereas diversification improves stability at the cost of some growth. The magnitude of these effects is larger for industrial structure than for fisheries, which is expected given the direct relevance of

	(1)	(2)	(3)	(1*)	(2*)	(3*)
NMP_{s,x_4}	0.024***	0.023***	0.024***	0.019***	0.018***	0.019**
	(0.007)	(0.007)	(0.007)	(0.006)	(0.006)	(0.006)

Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 2.7: Test results for normalized MP_{s,x_4}

industrial composition to the local economy, while fisheries diversification primarily influences the fishing sector with second-order effects on the broader economy.

The interaction-term coefficients between inputs, α_{kl}, capture complementary and substitution relationships in the production process (Morrison-Paul et al., 2000; Cuesta et al., 2009). Following Morrison-Paul et al. (2000)[5], a positive coefficient $\alpha_{kl} > 0$ indicates substitution, where the marginal product of input k decreases as input l increases (by symmetry, $\alpha_{kl} = \alpha_{lk}$). Conversely, a negative coefficient $\alpha_{kl} < 0$ indicates complementarity, where the marginal product of input k increases with input l.

A key result is the significance of the interaction between fisheries and industrial specialization. Estimates for α_{34} indicate substitution between these inputs: statistically significant in the EHDF and marginally significant in the HDF ($p < 0.13$ for HDF (1)–(3)). This suggests that fisheries and industrial specialization can substitute for each other, with implications for economic development in Alaskan fishing communities.

If physical or economic constraints limit specialization in the local economy, communities may instead specialize in fisheries to achieve growth. Conversely, if fisheries specialization is constrained by ecological or market factors, communities may specialize within a particular industrial sector. Thus, when specialization in one domain is constrained by geography or resource availability, communities may compensate by specializing in the other.

From the diversification perspective (the inverse of specialization), this implies complementarity between specialization and diversification across fisheries and the local fishing economy. Specializing in one domain while diversifying the other can support sustainable growth with greater stability. Balancing specialization and diversification across industrial sectors and fisheries creates a complementary dynamic that fosters stability without sacrificing growth.[6]

[5] Unlike most of the distance-function literature, which uses a negative sign on the left-hand side for normalization, Morrison-Paul et al. (2000) use a positive sign to simplify interpretation of production relationships. By flipping the sign of our estimated coefficients, we provide an interpretation consistent with Morrison-Paul et al. (2000). Similarly, Cuesta et al. (2009) use a negative sign but provide the same interpretation.

[6] As a robustness check, we re-estimate this complementary relationship using a direct diversification measure and assess which direction of complementarity is clearer (i.e., industrial diversification & fisheries specialization vs. industrial specialization & fisheries diversification).

In the SFA inefficiency equations that allow for heteroskedastic inefficiency with respect to fisheries-related community factors, the fisheries-permit share (a fisheries concentration index) shows a negative sign consistently, which aligns with challenges faced by small fishing communities with limited access to fishing-support services (Lavoie and Himes-Cornell, 2019). Survey and network evidence indicate that remote Alaskan fishing communities rely heavily on larger hub communities (e.g., Anchorage, Homer, Wrangell, Fairbanks, Ketchikan, Sitka, Kodiak) for infrastructure and services; this reliance is associated with higher fishing costs and reduced adaptive capacity to climate-related shocks.

Lastly, the fishing-dependency variable is positively associated with inefficiency, indicating that higher dependence on fishing coincides with greater distance from the production frontier. This association is consistent with lower economic growth and greater instability, underscoring that heavy reliance on natural resources may hinder efficient growth in fishing communities.

2.5.1 Robustness check 1: Testing endogeneity of input variables

Our objective is to develop a comprehensive understanding of how industrial and fisheries diversification shape growth and stability in local fishing economies, and how they interact with other growth factors in Alaska's fishing communities, rather than to identify variable-specific causal effects. Nevertheless, endogeneity remains a concern that could complicate interpretation.

We have argued that our HDF estimates are relatively robust to endogeneity, particularly simultaneity. Even so, unobserved community characteristics may influence input variables and economic instability, and the limited literature on community growth determinants, together with scarce community-level data, leaves residual uncertainty.

Empirical studies often mitigate endogeneity by incorporating additional data (e.g., input prices) or by using instrumental variables (IV) within SFA (Dalheimer et al., 2024; Sauer and Latacz-Lohmann, 2015; Atkinson et al., 2003). Applying IV methods in SFA is challenging: valid community-level instruments are difficult to find; even with valid instruments, our setting faces (1) a modest cross-sectional sample (few communities), (2) substantial consumption of degrees of freedom due to numerous parameters, and (3) potential weak-instrument problems, particularly in the translog HDF. As Amsler et al. (2016) note, addressing endogeneity in a translog specification with multiple endogenous variables requires instruments for nonlinearly transformed terms (squares and interactions), which may be weak; in such cases, two-stage estimators (e.g., 2SLS) can be statistically inefficient.

A potential alternative is the True Fixed-Effects (TFE) model in SFA (Greene, 2005a,b), which can address omitted time-invariant community factors. However, TFE estimation suffers from the incidental-parameter problem in panels with many communities (N large) and a fixed, finite time dimension (T small).

As a robustness check, we therefore test for potential endogeneity using the Bayesian Mundlak–Chamberlain device (MCD), which models unobserved heterogeneity and helps validate our results. Following Dalheimer et al. (2024), who apply an HDF framework with comparable inputs, desirable outputs, and undesirable outputs, we implement the Bayesian MCD. The device originates with Mundlak (1978) and Chamberlain (1982) and is extended to a Bayesian framework by Griffiths and Hajargasht (2016); it is well suited to SFA contexts.

The Mundlak–Chamberlain device (MCD) addresses endogeneity in panel models by modeling the correlation between unit effects and observed regressors. Like fixed effects, it mitigates bias from omitted variables correlated with both the dependent variable and the inputs. Relative to fixed effects, MCD estimation (1) preserves between variation, improving efficiency; (2) flexibly accommodates nonlinear models while capturing unobserved heterogeneity; and (3) avoids the incidental-parameter problem because it does not estimate unit-specific intercepts (Wooldridge, 2010).

Building on this idea, Griffiths and Hajargasht (2016) propose a Bayesian test for input endogeneity in SFA, subsequently applied by Dalheimer et al. (2024). The model is:

$$\ln y_{it} = \ln f(X_{it}, s_{it}; \beta) - u_i + v_{it}, \tag{2.17}$$

$$H(u_i) = \bar{x_i}'\delta + \zeta_{it}, \tag{2.18}$$

where $H(u_i) = \ln(u_i)$ and $f(\cdot)$ is a Cobb–Douglas distance function.[7] The inefficiency specification $H(\cdot)$ includes time-averaged values $\bar{x}_i$ of industrial (x_3) and fisheries specialization (x_4) to test for endogeneity. Unlike frequentist confidence intervals, Bayesian credible intervals admit a direct probabilistic interpretation of parameter uncertainty given the data and priors.

Statistically significant posterior means in the inefficiency equation within a given credible interval (e.g., 95%) indicate potential endogeneity (Griffiths and Hajargasht, 2016; Dalheimer et al., 2024). Conversely, insignificance—evidenced by a large posterior standard deviation relative to the mean—implies weak evi-

[7] Following Griffiths and Hajargasht (2016) and Dalheimer et al. (2024), we adopt Cobb–Douglas for computational efficiency in sampling and convergence.

dence of endogeneity. In Bayesian terms, the credible interval for the correlation parameters δ quantifies this uncertainty; wide intervals centered near zero suggest weak correlation.

We adopt the priors in Griffiths and Hajargasht (2016) for the frontier parameters β and the inefficiency parameters δ. To form the panel, we split the study period into two blocks, 2000–2008 and 2009–2016, yielding a structure that balances long-run trends and data availability. After excluding communities with insufficient time observations, the sample comprises $N = 165$ communities observed over $T = 2$ periods ($NT = 330$).[8] Table 2.8 reports the Bayesian MCD estimates.

$\ln y_1$		**Posterior Mean**	**Posterior SD**
Frontier			
Const.		0.233	0.208
$\ln x_1$	(Population)	2.770**	0.642
$\ln x_2$	(Wage Income per Capita)	0.413	0.457
$\ln x_3$	(Ind. Specialization)	1.707**	0.647
$\ln x_4$	(Fish. Specialization)	-0.759	0.430
$\ln s_1$	(Instability)	15.790**	6.829
Inefficiency			
Const.		-1.821**	0.918
$\overline{\ln x_3}$		0.008	1.409
$\overline{\ln x_4}$		0.006	1.422
Sample Observations		330	
** Significant at the 95% credible interval level.			

Table 2.8: Results from Bayesian Mundlak–Chamberlain device estimation.

We find no strong evidence of endogeneity for industrial or fisheries specialization, suggesting that endogeneity is not a major concern in our HDF model. The Bayesian estimates closely align with the HDF results: population (x_1) and industrial specialization (x_3) have positive effects on economic growth, while wage per capita (x_2) is insignificant. Fisheries specialization (x_4) is not significant, which may reflect the short panel length or specification differences (e.g., inclusion of quadratic terms). Instability shows a significant positive association with employment growth, reinforcing the link between higher growth and

[8] We set $T = 2$ periods (2000–2008 and 2009–2016) to ensure that each sub-period contains a sufficient number of annual observations (approximately $t = 8$) for reliably estimating the long-run growth trend and its associated instability via the geometric standard deviation. Splitting the sample into $T \geq 3$ periods reduces the number of within-period observations, which compromises the precision of sub-period growth and instability estimates. Preliminary analyses confirmed that finer temporal partitions yielded unstable estimates of these measures, supporting our choice of $T = 2$.

greater instability. Overall, the Bayesian MCD results support our main findings.

2.5.2 Robustness check 2: Estimation with direct diversification measures

In our main estimation, we invert industrial and fisheries diversification to interpret specialization and to maintain the non-increasing property of inputs, consistent with conventional input–output relationships in the HDF framework. To confirm robustness, we also estimate the model using direct diversification measures. This provides direct evidence of the complementary relationship between fisheries diversification (or specialization) and economic specialization (or diversification), rather than relying solely on indirect inference from specialization in the main analysis. Switching from fisheries specialization to its direct diversification counterpart yields coefficients with the same significance but opposite signs.

Complete results appear in Table C1 in Appendix C, which reports HDF estimates using the direct diversification measure (fisheries diversification, the inverse of specialization). The results align with our primary findings: the signs and significance of fisheries diversification and its interaction terms are consistent, with coefficients of similar magnitude but opposite signs. In addition, the interaction between fisheries diversification and economic specialization reaffirms the complementary relationship between diversification and specialization.

2.5.3 Robustness check 3: Estimation with alternative diversification measures using the Herfindahl–Hirschman index

We examine the sensitivity of our results to alternative diversification measures. In place of our default measure (based on the Shannon diversity index), we construct an alternative index using the inverse Herfindahl–Hirschman Index (HHI), also known as the Simpson diversity index, following Kluge (2018). Like the Shannon index, this measure is lower-bounded by 1. We define the inverse-HHI diversification index, Div_c^{HHI}, using sector shares P_{ict} (as in Equation (2.2)) as:

$$Div_c^{HHI} = \frac{1}{T}\sum_{t=1}^{T}\left(\frac{1}{\sum_{i=1}^{N} P_{ict}^2}\right). \tag{2.19}$$

Table D1 in Appendix D reports estimates using this alternative index (converted to a specialization measure for estimation). Results are qualitatively consistent with the baseline that uses the Shannon index,

differing only slightly in magnitude because the two measures are on different scales.[9]

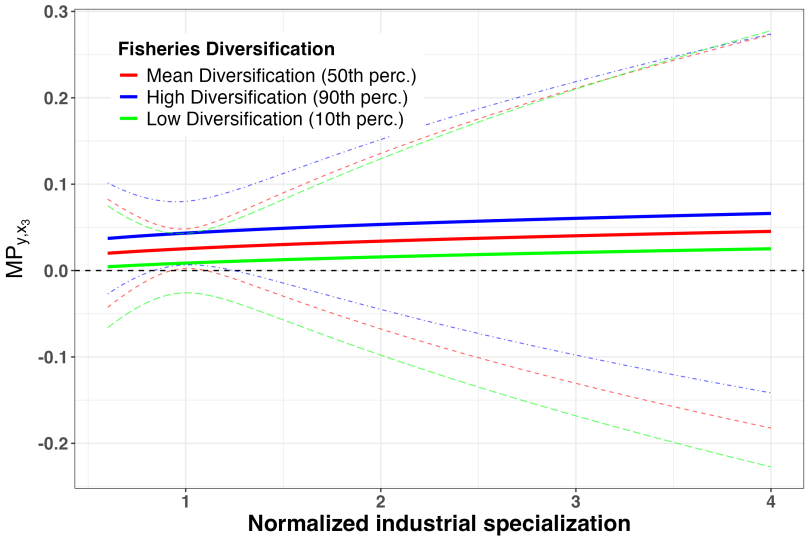
2.6 Discussion

The growth and stability of Alaska's local fishing economies are critical for sustaining fishing communities and fishers' livelihoods as climate-induced changes reshape ocean and coastal systems. Prior research on fisheries diversification has largely focused on stabilizing revenue within the fishing sector, often extrapolating broader community effects without direct empirical evidence. We address this gap by providing empirical evidence that fisheries diversification—alongside industrial diversification within the local economy—shapes the growth and stability of local fishing economies. In particular, both fisheries and industrial diversification are associated with greater stability in the local economies of fishing communities.

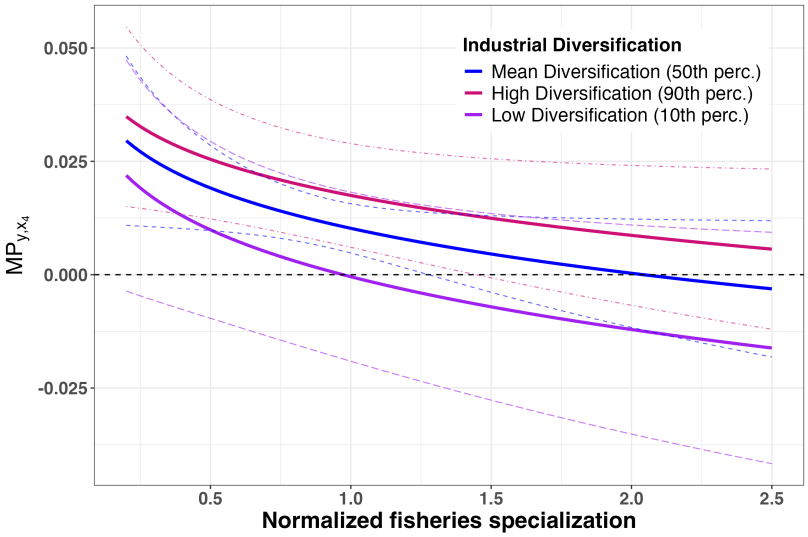
Our research provides the first empirical evidence that industrial and fisheries diversification are interlinked with each other and with broader community growth factors in shaping local economic growth. We find potential synergistic effects, where fisheries specialization and industrial diversification reinforce one another. Figure 2.4 shows the marginal product of each specialization across varying diversification levels in the complementary sector.

Figure 2.4a shows that the marginal product of industrial specialization increases monotonically with little evidence of diminishing returns. Greater fisheries diversification shifts this curve upward, indicating synergy; however, statistical precision is weaker for the industrial specialization estimates, as reflected in wide 95% confidence intervals. For fisheries specialization (Figure 2.4b), the marginal product exhibits diminishing returns: at very high levels of specialization, additional specialization can reduce local economic growth. As with industrial specialization, greater industrial diversification raises the marginal product curve and moderates the rate of diminishing returns. Moreover, the synergistic direction is statistically clearer for fisheries specialization than the complementary effect associated with industrial specialization. Overall, our results—supported by both estimations and visualizations—demonstrate that economic growth and stability in local fishing economies, driven by industrial and fisheries diversification, exhibit heterogeneous responses significantly influenced by the interaction with diversification levels in the complementary counterpart, whether in fisheries or the industrial structure of communities.

[9] A simple correlation analysis shows strong agreement between the two indices: the correlation coefficient is 0.97 for both industrial and fisheries diversification. These findings confirm that our main results are robust to the choice of diversification measure.



(a) Industrial specialization on growth



(b) Fisheries specialization on growth

Figure 2.4: Marginal products of industrial and fisheries specialization with complementary effects. Dashed lines denote 95% confidence intervals computed via the delta method (Appendix E). Estimates are based on specification (3) in Table 2.2 and are evaluated at the sample means of the other covariates. We use the HDF specification for interpretability because the EHDF case normalizes input x.

Previous studies have focused on the stabilizing effects of diversification while overlooking its trade-off with long-term growth. This narrow focus risks neglecting the vital role of economic growth in sustaining local fishing economies. Our findings show that, on average, Alaskan fishing communities have experienced a zero-growth steady state in employment, suggesting that achieving stability at the expense of growth may undermine sustainability. Therefore, a balanced approach that fosters both economic stability and growth—whether through fisheries diversification or a specialized industrial structure—is essential for long-

term viability.

Sethi et al. (2014) noted that geographical factors may constrain fisheries diversification, thus limiting its potential to drive economic stability or growth in certain contexts. In line with this, our findings suggest that economic stability can still be achieved through industrial structural improvements, provided that geographic constraints do not also restrict industrial reform. Communities can pursue industrial specialization for direct growth or diversification to enhance growth—particularly in support of specialized fisheries through the complementary effects shown in Figure 2.4. This dual-track strategy enables fishing communities to achieve balanced economic growth and stability by adapting fisheries and industrial structures to their specific community contexts, minimizing trade-offs between economic growth and stability.

Figure 2.4 further reveals that the complementary effects between fisheries specialization and industrial diversification are more evident than those between industrial specialization and fisheries diversification. While specializing in high-yield fisheries can drive revenue growth and generate spillover effects, it also heightens vulnerability to market and ecological shocks. A diversified industrial base—including processing, logistics, tourism, and trade—provides essential complementary services and alternative income sources during fisheries downturns. Consequently, fisheries-specialized communities may need to enhance resilience against a rapidly changing fishing environment by pursuing industrial diversification.

Overall, our research offers several key policy insights. First, fisheries policies alone are insufficient; they must be integrated with broader industrial development strategies to sustain local fishing economies. Second, place-based interventions that promote cross-sector investments and industrial diversification are essential for balancing the trade-off between economic growth and stability. Finally, tailored workforce development and community-specific strategies are crucial, as the optimal transformation of industrial structure depends on local conditions. Collectively, these insights highlight the need for a comprehensive, locally attuned policy framework to foster resilient and sustainable economic growth in Alaska's fishing communities.

Appendices

Appendix

A Estimation procedure details in stochastic frontier analysis

Our SFA employs a simultaneous estimation of the efficient frontier and inefficiency equations using Maximum Likelihood Estimation (MLE), avoiding the biases of the traditional two-step approach. To ensure reliable results, we address three critical factors: (i) the distributional assumption for the inefficiency term, (ii) initial parameter selection, and (iii) the optimization algorithm. These steps are essential to prevent convergence issues, avoid local maxima, and produce robust parameter estimates.

To determine initial values for maximizing the likelihood function, we followed the approach outlined by Kumbhakar et al. (2015). Specifically, we began by performing Ordinary Least Squares (OLS) estimations on both the frontier and inefficiency equations. The resulting coefficients were then used as initial values for the subsequent SFA process. This method provides a robust foundation for the optimization process and improves the reliability and precision of our estimations, particularly in cases where the default initial values generated by STATA may not perform adequately.

After evaluating both half-normal and truncated-normal distributions and finding non-significant results for the mean locus of the truncated-normal distribution, we opted for the half-normal distribution as proposed by Aigner et al. (1977), which is generally employed in the HDF literature. This choice was made for its simplicity and effectiveness in estimating the likelihood function, attributed to its single-parameter characteristic, unlike the two-parameter (mean locus and standard deviation) structure of the truncated-normal distribution.[10]

The selection of an optimization algorithm for the likelihood function significantly influences the stability and convergence of the likelihood function. Within our estimation environment, specifically STATA, we utilized four different optimization algorithms.[11] After experimenting with these algorithms, we identified those that exhibited stable convergence of the likelihood function, along with the highest likelihood value upon stable convergence. This approach ensures that our choice of optimization algorithm not only promotes efficient convergence but also optimizes the accuracy and reliability of our estimation results.

[10] Exponential and gamma distributions were also considered but the half-normal and truncated-normal distributions showed superior convergence in our preliminary analysis, leading us to focus on these distributions.

[11] Newton-Raphson, Broyden–Fletcher–Goldfarb–Shanno (BFGS), Davidon–Fletcher–Powell (DFP), and Berndt–Hall–Hall–Hausman (BHHH).

B Measuring economic growth rate and economic instability

In our study, we utilized the geometric growth rate of employment as our measure for average growth. This metric is often employed to present a more accurate long-term performance of a portfolio. It is based on the principle that performance in one period affects subsequent periods. Conversely, the arithmetic average may not accurately measure long-term growth when economic volatility is prevalent.

We adopt the methodologies utilized by Lande (1994), Chandra (2002), and Hafner (2019). To elucidate this approach, we begin with a continuous-time exponential model for employment growth rate, which is expressed as:

$$Emp_t = Emp_0 \cdot e^{\beta_1 \cdot t}$$

Here, Emp_0 represents the initial employment at time $t = 0$, and β_1 is the instantaneous employment growth rate. By logarithmically transforming both sides of the equation, we obtain:

$$\log(Emp_t) = \beta_0 + \beta_1 \cdot t$$

where $\beta_0 = \log(Emp_0)$, treated as a constant. This log-transformed equation, after including an error term ϵ_t, is estimated using Ordinary Least Squares (OLS) as follows:

$$\log(Emp_t) = \beta_0 + \beta_1 \cdot t + \epsilon_t$$

With the estimated coefficient of the time trend, $\hat{\beta}_1$, interpreted as the average instantaneous growth rate of employment, we difference the estimated equation from $t - 1$ to t:

$$\log(Emp_t) - \log(Emp_{t-1}) = \hat{\beta}_1$$

By exponentiating the above equation and recognizing that the discrete-time version of exponential growth corresponds to the geometric growth rate, we establish the following relationship:

$$(1 + \mu) = \exp(\hat{\beta}_1)$$

where μ represents the average annual geometric growth rate of employment for a community. Since the HDF employs a log-transformation to ensure all values remain positive, we use the average annual geometric growth *return*, $y_c = \exp(\beta_1, c)$, where the subscript c denotes each community.

Subsequently, instability is quantified as the standard deviation of residuals from an OLS estimation based on the previous log-time trend equation. The fitted regression line on the time trend captures the long-term (geometric) return; thus, larger deviations from this fitted line, whether positive or negative, are interpreted as economic instability associated with year-to-year economic growth. The sum of squared residuals is computed as follows:

$$\sum_{t=1}^{T} \left(\ln Emp_t - \ln \hat{Emp}_t\right)^2 = \sum_{t=1}^{T} \left(\ln \frac{Emp_t}{\hat{Emp}_t}\right)^2$$

Along with the geometric mean, the formula for geometric standard deviation (GSD), σ_g, is expressed as an exponentiated arithmetic mean of the logged differences between some values A_n and their geometric mean μ_g over a number of observations N as follows:

$$\text{GSD} = \sigma_g = \exp\left(\sqrt{\frac{1}{N}\sum_{n=1}^{N}\left(\ln \frac{A_n}{\mu_g}\right)^2}\right).$$

Following this formula, the equivalent growth instability of fishing community c, as the geometric standard deviation of the long-term growth trend, can be derived using the fitted mean regression line on time trend ($\ln(\hat{Emp}_t) = \hat{\beta}_0 + \hat{\beta}_1 \cdot t$), as follows. The economic instability of community c, represented as the geometric standard deviation of economic growth, is given by:

$$\sigma_c = s_c = \exp\left(\sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(\ln \frac{Emp_t}{\hat{Emp}_t}\right)^2}\right).$$

C Robustness check using direct diversification measure

	(1)	(2)	(3)	(1*)	(2*)	(3*)
	HDF			**EHDF**		
Population (α_1)	-0.017***	-0.018***	-0.019***	-0.018***	-0.018***	-0.020***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Wage/Capita Inc. (α_2)	0.002	0.001	-0.000	0.001	0.000	-0.001
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Ind. Specialization (α_3)	-0.014*	-0.015*	-0.018**	-0.012	-0.013*	-0.016**
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Fish. Diversification (α_4)	0.007***	0.007***	0.007***	0.006***	0.006***	0.006***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Population sq. (α_{11})	0.008***	0.008***	0.009***	0.008***	0.008***	0.008***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Wage/Capita Inc. sq. (α_{22})	-0.021	-0.023*	-0.027**	-0.020	-0.022	-0.026*
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.013)
Ind. Specialization sq. (α_{33})	-0.002	0.001	-0.005	-0.018	-0.016	-0.021
	(0.045)	(0.045)	(0.045)	(0.044)	(0.044)	(0.044)
Fish. Diversification sq. (α_{44})	0.010**	0.009*	0.010**	0.010**	0.009*	0.010**
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Population * Wage/Capita Inc. (α_{12})	-0.006	-0.006	-0.007	-0.005	-0.004	-0.005
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Population * Ind. Special. (α_{13})	0.005	0.005	0.001	0.002	0.002	-0.002
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Population * Fish. Div. (α_{14})	-0.008***	-0.008***	-0.008***	-0.008***	-0.007***	-0.008***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Wage/Capita Inc. * Ind. Special. (α_{23})	-0.017	-0.018	-0.023	-0.023	-0.025	-0.030
	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)
Wage/Capita Inc. * Fish. Div. (α_{24})	0.003	0.003	0.004	0.001	0.001	0.002
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Econ Special. * Fish. Div. (α_{34})	-0.015	-0.015	-0.013	-0.016*	-0.016*	-0.014
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Econ. Instability (δ_1)	-0.302***	-0.296***	-0.298***	-0.321***	-0.314***	-0.317***
	(0.040)	(0.040)	(0.039)	(0.040)	(0.039)	(0.039)
Econ. Instability sq (δ_{11})	0.281	0.420	0.370	0.617	0.743	0.664
	(0.912)	(0.929)	(0.942)	(0.860)	(0.875)	(0.886)
Population * Econ. Instability (η_{11})	-0.017	-0.014	-0.016	-0.007	-0.010	-0.007
	(0.041)	(0.041)	(0.041)	(0.040)	(0.041)	(0.040)
Wage/Capita Inc. * Econ. Instability (η_{21})	-0.416***	-0.402***	-0.425***	-0.349***	-0.336***	-0.364***
	(0.116)	(0.115)	(0.114)	(0.114)	(0.113)	(0.111)
Ind. Special. * Econ. Instability (η_{31})	-0.073	-0.083	-0.116	-0.014	-0.022	-0.059
	(0.167)	(0.167)	(0.167)	(0.162)	(0.163)	(0.163)
Fish. Diversification * Econ. Instability (η_{41})	0.064	0.053	0.055	0.038	0.025	0.028
	(0.050)	(0.051)	(0.051)	(0.049)	(0.050)	(0.050)
Inefficiency						
Broader Permit Share		-0.720	-1.004		-0.717	-1.024*
		(0.734)	(0.644)		(0.703)	(0.618)
Fishing Earnings to Wage Income Ratio			0.051**			0.052**
			(0.025)			(0.023)
Observations	177	177	177	177	177	177

Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table C1: Estimation results for robustness check: using fisheries diversification

D Robustness check using alternative diversification measure with inverse HHI index

	(1)	(2)	(3)	(1*)	(2*)	(3*)
	HDF			**EHDF**		
Population (α_1)	-0.017***	-0.018***	-0.018***	-0.018***	-0.018***	-0.020***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Wage/Capita Inc. (α_2)	0.002	0.001	0.000	0.001	0.000	-0.001
	(0.004)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)
Ind. Specialization (α_3)	-0.014*	-0.015**	-0.017**	-0.012	-0.013*	-0.016**
	(0.007)	(0.007)	(0.007)	(0.008)	(0.008)	(0.008)
Fish. Diversification (α_4)	-0.009***	-0.008***	-0.009***	-0.006***	-0.006***	-0.006***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Population sq. (α_{11})	0.008***	0.008***	0.008***	0.008***	0.008***	0.008***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
Wage/Capita Inc. sq. (α_{22})	-0.020	-0.023*	-0.026*	-0.020	-0.022	-0.026*
	(0.014)	(0.014)	(0.013)	(0.014)	(0.014)	(0.013)
Ind. Specialization sq. (α_{33})	0.015	0.018	0.017	-0.018	-0.016	-0.021
	(0.047)	(0.047)	(0.047)	(0.044)	(0.044)	(0.044)
Fish. Diversification sq. (α_{44})	0.012**	0.011*	0.011*	0.010**	0.009*	0.010**
	(0.006)	(0.006)	(0.006)	(0.005)	(0.005)	(0.005)
Population * Wage/Capita Inc. (α_{12})	-0.005	-0.005	-0.005	-0.005	-0.004	-0.005
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Population * Ind. Special. (α_{13})	0.007	0.008	0.005	0.002	0.002	-0.002
	(0.008)	(0.008)	(0.008)	(0.009)	(0.009)	(0.009)
Population * Fish. Div. (α_{14})	0.008***	0.008***	0.009***	0.008***	0.007***	0.008***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Wage/Capita Inc. * Ind. Special. (α_{23})	-0.009	-0.011	-0.013	-0.023	-0.025	-0.030
	(0.019)	(0.019)	(0.019)	(0.018)	(0.019)	(0.019)
Wage/Capita Inc. * Fish. Div. (α_{24})	-0.004	-0.004	-0.005	-0.001	-0.001	0.002
	(0.006)	(0.006)	(0.006)	(0.005)	(0.005)	(0.005)
Econ Special. * Fish. Div. (α_{34})	0.018	0.018	0.016	-0.016*	-0.016*	-0.014
	(0.011)	(0.011)	(0.011)	(0.010)	(0.010)	(0.010)
Econ. Instability (δ_1)	-0.307***	-0.301***	-0.303***	-0.321***	-0.314***	-0.317***
	(0.040)	(0.040)	(0.039)	(0.040)	(0.039)	(0.039)
Econ. Instability sq. (δ_{11})	0.138	0.298	0.231	0.617	0.743	0.664
	(0.897)	(0.918)	(0.935)	(0.860)	(0.875)	(0.886)
Population * Econ. Instability (η_{11})	-0.017	-0.015	-0.016	-0.007	-0.010	-0.007
	(0.038)	(0.038)	(0.038)	(0.040)	(0.041)	(0.040)
Wage/Capita Inc. * Econ. Instability (η_{21})	-0.454***	-0.440***	-0.459***	-0.349***	-0.336***	-0.364***
	(0.112)	(0.111)	(0.111)	(0.114)	(0.113)	(0.111)
Ind. Special. * Econ. Instability (η_{31})	-0.084	-0.098	-0.121	-0.014	-0.022	-0.059
	(0.166)	(0.167)	(0.168)	(0.162)	(0.163)	(0.163)
Fish. Diversification * Econ. Instability (η_{41})	-0.078	-0.065	-0.066	-0.036	-0.025	-0.028
	(0.055)	(0.056)	(0.057)	(0.049)	(0.050)	(0.050)
Inefficiency						
Broader Permit Share		-0.728	-0.950		-0.717	-1.024*
		(0.695)	(0.633)		(0.703)	(0.618)
Fishing Earnings to Wage Income Ratio			0.044*			0.052**
			(0.026)			(0.023)
Observations	177	177	177	177	177	177

Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table D1: Estimation results for robustness check: using an alternative diversification measure (inverse HHI)

E Constructing 95% confidence intervals for MP_{y,x_3,x_4}

We outline the detailed procedure for constructing a confidence interval for the marginal product (MP) of input to output, calculated as the ratio of two distance elasticities. These elasticities are derived from differentiating the hyperbolic distance function (HDF). Because the elasticities, except for the variable of interest, are evaluated at their mean values, the remaining logged variables become zero after normalization by geometric means. This approach allows us to examine how the MP changes with respect to a single variable of interest—industrial or fisheries specialization—while holding all other factors constant (*ceteris paribus*). For illustrative purposes, we focus on the procedure for industrial specialization. The distance elasticities for industrial specialization and employment growth are presented below.

$$\epsilon_{D,x3} = \alpha_3 + \alpha_{33} \log(x3)$$

$$\epsilon_{D,y} = 1 + \delta_1 + \eta_{31} \log(x3)$$

Our goal is to compute the variance of $MP_{y,x3}$, which is the ratio of these elasticities:

$$MP = R(\epsilon_{D,x3}, \epsilon_{D,y}) = -\frac{\epsilon_{D,x3}}{\epsilon_{D,y}}$$

To find the variance of R, we use the Delta Method, which provides an approximation for the variance of a function of random variables. The Delta Method states that for a function $g(\boldsymbol{\theta})$ of random variables $\boldsymbol{\theta}$, the variance of $g(\boldsymbol{\theta})$ can be approximated as:

$$\text{Var}[g(\boldsymbol{\theta})] \approx \nabla g(\boldsymbol{\theta})^\top \boldsymbol{\Sigma} \nabla g(\boldsymbol{\theta})$$

where $\nabla g(\boldsymbol{\theta})$ is the gradient vector of partial derivatives of g with respect to $\boldsymbol{\theta}$, and $\boldsymbol{\Sigma}$ is the covariance matrix of $\boldsymbol{\theta}$.

We begin by computing the partial derivatives of R with respect to $\epsilon_{D,x3}$ and $\epsilon_{D,y}$:

$$\frac{\partial R}{\partial \epsilon_{D,x3}} = -\frac{1}{\epsilon_{D,y}}$$

$$\frac{\partial R}{\partial \epsilon_{D,y}} = -\left(-\frac{\epsilon_{D,x3}}{\epsilon_{D,y}^2}\right) = \frac{\epsilon_{D,x3}}{\epsilon_{D,y}^2}$$

Applying the Delta Method, the variance of R is approximated as:

$$\text{Var}(R) \approx \left(\frac{\partial R}{\partial \epsilon_{D,x3}}\right)^2 \text{Var}(\epsilon_{D,x3}) + \left(\frac{\partial R}{\partial \epsilon_{D,y}}\right)^2 \text{Var}(\epsilon_{D,y}) + 2\left(\frac{\partial R}{\partial \epsilon_{D,x3}}\right)\left(\frac{\partial R}{\partial \epsilon_{D,y}}\right) \text{Cov}(\epsilon_{D,x3}, \epsilon_{D,y})$$

Substituting the expressions for the partial derivatives:

$$\text{Var}(R) \approx \left(-\frac{1}{\epsilon_{D,y}}\right)^2 \text{Var}(\epsilon_{D,x3}) + \left(\frac{\epsilon_{D,x3}}{\epsilon_{D,y}^2}\right)^2 \text{Var}(\epsilon_{D,y}) + 2\left(-\frac{1}{\epsilon_{D,y}}\right)\left(\frac{\epsilon_{D,x3}}{\epsilon_{D,y}^2}\right) \text{Cov}(\epsilon_{D,x3}, \epsilon_{D,y})$$

Simplifying the expression:

$$\text{Var}(R) \approx \frac{\text{Var}(\epsilon_{D,x3})}{\epsilon_{D,y}^2} + \frac{\epsilon_{D,x3}^2 \, \text{Var}(\epsilon_{D,y})}{\epsilon_{D,y}^4} - \frac{2\epsilon_{D,x3} \, \text{Cov}(\epsilon_{D,x3}, \epsilon_{D,y})}{\epsilon_{D,y}^3}$$

Next, we compute the variances of $\epsilon_{D,x3}$ and $\epsilon_{D,y}$. Since $\epsilon_{D,x3}$ is a linear combination of α_3 and α_{33}, its variance is:

$$\text{Var}(\epsilon_{D,x3}) = \text{Var}(\alpha_3) + [\log(x3)]^2 \, \text{Var}(\alpha_{33}) + 2\log(x3) \, \text{Cov}(\alpha_3, \alpha_{33})$$

However, as the covariance $\text{Cov}(\alpha_3, \alpha_{33})$ is negligible in our estimations, we simplify:

$$\text{Var}(\epsilon_{D,x3}) \approx \text{Var}(\alpha_3) + [\log(x3)]^2 \, \text{Var}(\alpha_{33})$$

Similarly, the variance of $\epsilon_{D,y}$ is:

$$\text{Var}(\epsilon_{D,y}) = \text{Var}(\delta_1) + [\log(x3)]^2 \, \text{Var}(\eta_{31}) + 2\log(x3) \, \text{Cov}(\delta_1, \eta_{31})$$

Assuming $\text{Cov}(\delta_1, \eta_{31})$ is negligible, we have:

$$\text{Var}(\epsilon_{D,y}) \approx \text{Var}(\delta_1) + [\log(x3)]^2 \, \text{Var}(\eta_{31})$$

The covariance between $\epsilon_{D,x3}$ and $\epsilon_{D,y}$ is given by:

$$\text{Cov}(\epsilon_{D,x3}, \epsilon_{D,y}) = \text{Cov}(\alpha_3 + \alpha_{33}\log(x3), \delta_1 + \eta_{31}\log(x3))$$

Expanding this covariance:

$$\text{Cov}(\epsilon_{D,x3}, \epsilon_{D,y}) = \text{Cov}(\alpha_3, \delta_1) + \log(x3)[\text{Cov}(\alpha_3, \eta_{31}) + \text{Cov}(\alpha_{33}, \delta_1)] + [\log(x3)]^2 \text{Cov}(\alpha_{33}, \eta_{31})$$

In our estimation, the covariance $\text{Cov}(\epsilon_{D,x3}, \epsilon_{D,y})$ consists of covariances between various parameters, $\alpha_3, \delta_1, \eta_{31}, \alpha_{33}$, which include a mix of negative and positive values, thus, the sum of the terms are close to zero. These terms are multiplied by $\log(x3)$ and $[\log(x3)]^2$, further reducing their already small magnitude. For simplicity and based on this justification, we assume the covariance between these two elasticities is negligible, effectively zero:

$$\text{Cov}(\epsilon_{D,x3}, \epsilon_{D,y}) \approx 0$$

With this assumption, the simplified variance of R, derived using the delta method, becomes:

$$\text{Var}(MP_{y,x_3}) \approx \frac{\text{Var}(\epsilon_{D,x3})}{\epsilon_{D,y}^2} + \frac{\epsilon_{D,x3}^2 \, \text{Var}(\epsilon_{D,y})}{\epsilon_{D,y}^4}$$

Once the variance is computed, the 95% confidence interval for the marginal product of x_3 is constructed using its root-squared value. For the marginal product related to fisheries specialization, an identical procedure is followed.

Chapter 3

Balancing Multiple Interventions for Dynamically Efficient Kelp Forest Restoration Under Marine Heatwave Uncertainty

Abstract

Kelp forests are declining globally, threatening key socio-ecological functions such as marine refugia, coastal productivity, and community livelihoods. Managers face difficult decisions about when, how, and how intensively to restore such systems under complex feedbacks and uncertainty. In Northern California, recent marine heatwaves (MHWs) and elevated kelp grazing by purple sea urchins have caused severe kelp loss and economic hardship from collapse of the red urchin fishery. We develop a dynamic bioeconomic model of kelp-urchin system that is the first to identify welfare-maximizing restoration actions that flexibly depend on the current state of the system. The model integrates (i) kelp-urchin feedback dynamics, (ii) two intensity-varying restoration tools, kelp outplanting and purple urchin removal, (iii) restoration costs and benefits, and (iv) stochastic MHW regimes with uncertain frequency. We find contrasting dynamics between restoration tools: outplanting is critical when kelp population is low and is scaled back once kelp recovers to moderate levels, while optimal urchin removal is aggressive across most states. Furthermore, we assess the sensitivity of baseline results to alternative natural and social system scenarios. Our results are only modestly sensitive to the biological and climate alternatives we consider, scaling in intuitive fashion. In contrast, welfare and policy shift substantially under the social and management alternatives assessed. For example, restricting the functional form of policy response to a commonly used heuristic results in surprisingly high losses of welfare as poor fit to the optimal policy for one action feeds back into the other.

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

3.1 Introduction

Kelp forests are vital to many coastal marine ecosystems. They offer 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, 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 urchin grazing pressure on living kelp, leading to 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. For example, in 2017, kelp loss led to an estimated 80% mortality in the red abalone population, resulting in closure of this recreational fishery in 2018 (Rogers-Bennett and Catton, 2019). Similarly, the red sea urchin commercial 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). This fishery previously generated $2-3 million in annual fishing revenues, benefiting urchin divers, local community members employed by seafood processing companies for urchin processing (extracting uni), and various other related jobs. When fisheries are disrupted, this can lead to the loss of other important but intangible assets, such as specialized professional knowledge for this dive fishery. Furthermore, indigenous communities in California have also been heavily impacted, given their historical reliance 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, accounting for both ecosystem services (e.g., carbon sequestration and water purification) and as well as commercial and recreational fisheries across different 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 biological model, the approach does not incorporate restoration costs, and the simplified management heuristic of employing a static rule for restoration intensity that is static over time does not provide the flexibility needed to address the reality of evolving restoration needs. Both studies assume deterministic restoration outcomes, setting aside the strong stochasticity inherent in kelp systems. While these analyses assess the efficacy of management options constrained in time and/or form, no existing study has solved for dynamically optimal kelp forest restoration that flexibly responds to changes in the state of the system.

A growing body of literature illustrates the importance of taking an ecosystem-based management approach, capturing feedback between multiple species, resource management, and stakeholder well-being. 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 highlight the role of habitat management in enhancing 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, with no existing treatment of kelp forest restoration from a degraded state.

We assess kelp restoration strategies that balance 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, which simultaneously warm and de-nutrify coastal waters, causing catastrophic losses of canopy-forming kelps (McPherson et al., 2021). Our approach captures two-way feedbacks between kelp and urchin species as well as between the biological state of the system and management choices. In the first stage of our analysis, we establish baseline results for optimal state-contingent restoration in this system with nonlinear kelp-urchin feedbacks under stochastic MHW uncertainty. Previewing our results, we find that optimal kelp forest restoration includes a dynamic mix of kelp outplanting and purple urchin removal. Outplanting is particularly important and intensive when kelp population is low, while purple urchin removal is broadly used across most system states to keep grazing pressure in check.

Because the set of core natural and social components to the model is substantial, in the second stage of our analysis, we focus on identifying where and how results are sensitive to alternative model scenarios. We consider two categories of scenarios. First, we assess key biological and climate alternatives, including degraded versus healthy biological starting states, MHW-frequency regimes, and kelp spore connectivity specifications. In our second category, we focus on social and management factors, namely the implications of constraining either fishing or restoration tool use (in type or functional form). This distinction proves informative. Our results are only modestly sensitive to the biological and climate alternatives we consider, scaling in intuitive fashion. In contrast, welfare and policy shift substantially under the key social and management alternatives assessed. For example, the scale of impact of restricting restoration tool use to a single type diverges substantially by tool selected, due to asymmetry in substitutability. Additionally, restricting the functional form of policy response to a commonly used heuristic results in surprisingly high

losses of welfare as poor fit to the optimal policy for one action feeds back into the other.

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

The dynamic bioeconomic model tracks the population 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 model of kelp and purple urchin dynamics from Arroyo-Esquivel et al. (2023). We extend the model in three ways. First, we add the red urchin population. Second, we add uncertainty to the model in the form of stochastic annual occurrence of MHWs and uncertainty about the true MHW regime (annual likelihood). 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 beliefs about the MHW regime).

3.2.1 Biological dynamics

Our annual time step starts in the fall season to align with 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 (CDFW, 2020; Ward et al., 2022). This definition of the time step also aligns with the data available for model parameterization, which has been primarily collected during the fall. The spatial unit for our model is a 360 m^2 representative patch, matching the survey plot size in our parameterization dataset.

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 kelp

stipes remaining after "active" grazing by urchins is given by[1]

$$k(X_t, P_t, R_t) = X_t \cdot \underbrace{\exp\Big(- \gamma_P \cdot g_P(X_t, P_t) - \gamma_R \cdot g_R(X_t, R_t)\Big)}_{\text{active grazing rate on adult kelp}}. \tag{3.1}$$

The coefficients γ_P and γ_R correspond to the rate of decline in kelp abundance per unit of active grazing intensity from purple and red urchins, respectively. The function $g_U(\cdot),\ U \in \{P, R\}$, defining the grazing intensity from either urchin population, is described in detail next.

Under a kelp-abundant environment, urchins subsist on drift kelp blades that detach from kelp stipes 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, 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 an 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 active grazing intensity of urchin population $U \in \{P, R\}$ is given by

$$g_U(X_t, U_t) = \frac{U_t}{\sigma_1^U \cdot X_t^2 + \sigma_2^U \cdot X_t + 1}, \tag{3.2}$$

where σ_1^U drives the kelp density at which each urchin grazing intensity reaches its maximum and σ_2^U regulates the magnitude of this grazing intensity for urchin species U. This functional response represents a uni-modal and concave active grazing intensity that peaks at $X^U_{\text{max graze}} \equiv \frac{1}{\sqrt{\sigma_1^U}}$, the kelp density at which grazing per unit of kelp is maximized for a given level of urchin species U.[3]

[1] Arroyo-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. This exponential formulation provides three advantages: (i) it ensures smoothness and differentiability across the state space, which is important for the optimization algorithm in identifying the solutions; (ii) it follows the standard approach in population ecology for modeling mortality processes; and (iii) it realistically reflects the grazing impacts that accumulate continuously rather than abrupt switching at a threshold.

[2] Kelp 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.

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

Kelp dynamics

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 υ per unit of adult kelp, the resulting level of viable adult kelp recruits per unit of existing adult kelp, $k(X_t, P_t, R_t)$, is

$$y(X_t, P_t, R_t) = \frac{\upsilon \cdot \exp\Big(- \mu_P \cdot g_P(X_t, P_t) - \mu_R \cdot g_R(X_t, R_t)\Big)}{1 + \phi \cdot k(X_t, P_t, R_t)}, \tag{3.3}$$

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.[4] In our baseline specification, the recruitment function does not include external recruitment from outside of the patch because (1) many patches in our study area are surrounded by persistent urchin barrens, making adjacent patches unlikely to serve as reliable sources of kelp recruitment; (2) from a conservative restoration perspective, we avoid assuming automatic kelp recovery once urchin densities fall; and (3) external recruitment would likely be paired with outflow of local recruits and our baseline model is equivalent to the neutral assumption that any gain and loss of recruits is equal. Nonetheless, we later examine an alternative scenario which includes external recruitment.

Kelp population dynamics are given by

$$X_{t+1} = x(X_t, P_t, R_t, A_{X,t}|Z_t) = \underbrace{\Big[k(X_t, P_t, R_t) \cdot \Big(y(X_t, P_t, R_t) + \delta_X)\Big) \Big]}_{\text{kelp recruitment \& adult survivorship}} \cdot \underbrace{(1 - \xi \cdot Z_t)}_{\text{MHW shock}} + \underbrace{\eta_X \cdot A_{X,t}}_{\substack{\text{management:} \\ \text{outplanting}}}, \tag{3.4}$$

which combines kelp recruitment, density-independent natural mortality of adult kelp, a stochastic MHW shock and the restoration intervention of kelp outplanting $A_{X,t}$. δ_X is the coefficient of natural annual survival of adult kelp and η_X is the marginal effect of juvenile kelp outplanting. Finally, the binary variable Z_t indicates the occurrence of a MHW, which leads to mortality of a fraction, ξ, of kelp, a process we

[4] When kelp density is sufficiently high, the total recruitment level from the stock $k(X_t, P_t, R_t)$ approaches $\frac{\upsilon}{\phi}$ in the absence of urchin spore consumption.

describe in detail next.[5]

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 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 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). We generalize this stochastic component to allow for a MHW frequency α_i from a range of MHW regimes, $i \in \{low, medium, high\}$. The *high* regime frequency is based on the historical record from 2010-2020 during which MHWs were observed at an elevated frequency. The *low* regime frequency is based on the prior record (1985-2009) and *medium* is given by the average of the two extremes. See Appendix C.2 for details.

Urchin dynamics

To incorporate energy gained from kelp grazing in our model of urchin reproduction dynamics, we distinguish between energy obtained from active grazing on adult kelp (addressed here) and passive grazing on drift kelp (discussed in the following section).[6] 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 (CDFW, 2020). As a result, purple urchins exert the initial active grazing pressure on adult kelp. According to the adult kelp grazing model in Equation (3.1), they consume an amount of kelp equal to $X_t \cdot (1 - \exp(-\gamma_P \cdot g_P(X_t, P_t)))$. Multiplying this consumption by the energy conversion efficiency ϵ_1^P, the resulting purple urchin reproduction is

$$p_X(X_t, P_t) = \epsilon_1^P \cdot X_t \cdot (1 - \exp(-\gamma_P \cdot g_P(X_t, P_t))). \tag{3.5}$$

[5] Arroyo-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.

[6] We assume that energy intake from microscopic spore kelp, as modeled in Equation (3.3), is negligible and thus does not contribute meaningfully to reproduction, following the approach of Arroyo-Esquivel et al. (2023).

Red urchin consumption of adult kelp after purple urchins leads to reproduction, given by

$$r_X(X_t, P_t, R_t) = \epsilon_1^R \cdot X_t \cdot \exp(-\gamma_P \cdot g_P(X_t, P_t)) \cdot (1 - \exp(-\gamma_R \cdot g_R(X_t, R_t))), \tag{3.6}$$

where ϵ_1^R represents the energy conversion factor for red urchins.

Next, we model urchin reproductive gains from passive grazing of drift kelp as a function of kelp abundance and urchin population. We assume that red urchins passively graze first, followed by purple urchins. Reproduction from drift kelp consumption for both species is given by

$$r_D(X_t, R_t) = \epsilon_2^R \cdot X_t \cdot (1 - \exp(-\tau_R \cdot R_t)), \tag{3.7}$$

$$p_D(X_t, P_t, R_t) = \epsilon_2^P \cdot X_t \cdot \exp(-\tau_R \cdot R_t) \cdot (1 - \exp(-\tau_P \cdot P_t)), \tag{3.8}$$

where ϵ_2^R and ϵ_2^P are effective energy conversion parameters for red and purple urchins, respectively, while τ_R and τ_P capture the strength of drift kelp grazing effects by each urchin species. The full derivation of these equations and the biological assumptions underpinning them are provided in Appendix A.1.

Given reproduction dynamics specified above, purple urchin dynamics are given by

$$p(X_t, P_t, R_t, A_{P,t}) = \underbrace{\delta_P \cdot P_t}_{\substack{\text{urchin} \\ \text{survival}}} + \underbrace{p_X(X_t, P_t) + p_D(X_t, P_t, R_t)}_{\text{urchin recruitment}} - \underbrace{\eta_P \cdot A_{P,t}}_{\substack{\text{management:} \\ \text{removal}}}, \tag{3.9}$$

where δ_P is the annual survival proportion of purple urchins, and η_P specifies 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[7]

$$P_{t+1} = \max[P_{\text{exo}}\ ,\ p(X_t, P_t, R_t, A_{P,t})], \tag{3.10}$$

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

Consistent with historical fishing practices, we assume that the allowable level of red urchin harvest is given by a fixed proportion of the commercially valuable size class of the population. Additionally, red urchins are only harvested when they possess commercially valuable gonads ("uni"), since red urchins in food-deprived environments lack sufficient gonadal quality for market purposes. Red urchin harvest is given by

$$H_t = h(X_t, R_t) = \lambda \cdot \omega \cdot R_t \cdot \mathbb{1}(X_t > X^P_{\text{max graze}}), \tag{3.11}$$

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 are commercially valuable coincides with the kelp population level at which active grazing intensity per purple urchin is maximized, i.e., $X^P_{\text{max graze}}$, as defined in Equation (3.2). This assumption aligns with observations that in purple urchin barrens—where the kelp population is low and direct grazing of kelp stipes predominates—red urchins typically lack commercially viable gonads due to insufficient nourishment from kelp as a food source (Smith et al., 2021; Harrold and Reed, 1985). The choice is also consistent with the fact that purple urchins dominate grazing in urchin barrens, a pattern commonly observed in Northern California.[8]

Incorporating the previously established harvest rule and reproductive functions in Equations (3.6) and (3.7), along with natural mortality and recruitment, red urchin population dynamics are given by

$$r(X_t, P_t, R_t) = \delta_R \cdot R_t + \Big(r_X(X_t, P_t, R_t) + r_D(X_t, R_t)\Big) - H_t, \tag{3.12}$$

[7] This 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.

[8] Our parameterized model reflects this dominance, with purple urchins primarily responsible for grazing in low-kelp environments (see Fig. A1 in Appendix A.1). In addition, the parameter values determining maximum grazing for both species are similar in our model, leading to comparable peak grazing intensities at nearly the same kelp population levels, i.e., $X^P_{\text{max graze}} \approx X^R_{\text{max graze}}$.

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

$$R_{t+1} = \max[R_{\text{exo}},\ r(X_t, P_t, R_t)], \tag{3.13}$$

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

3.2.2 Management components

Building on the restoration management actions introduced above—kelp outplanting (A_X) and purple urchin removal (A_P)—we now specify the cost structure of these efforts, as well as the benefits derived from kelp, including its direct value and increased red urchin harvest. 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:

$$c_j(A_{j,t}) = \theta_j \cdot A_{j,t}. \tag{3.14}$$

When red urchin harvest takes place, a fixed cost is incurred, independent of density and harvest level:

$$c_H(H_t) = \theta_H \cdot \mathbb{1}(H_t \neq 0). \tag{3.15}$$

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., diving equipment provisioning and seafood processing). Second, we assume that standing kelp population provides direct ecological and use value, modeled as a linear function $\kappa_X \cdot X_t$. 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.

The reward function in period t is given by the net benefits:

$$\pi(H_t, X_t, A_{X,t}, A_{P,t}) = \underbrace{\kappa_R \cdot (\rho \cdot H_t - c_H(H_t))}_{\text{red urchin harvest profit and multiplier effects}} + \underbrace{\kappa_X \cdot X_t}_{\text{direct kelp value}} - \underbrace{c_X(A_{X,t})}_{\text{kelp outplanting cost}} - \underbrace{c_P(A_{P,t})}_{\text{purple urchin removal cost}}. \tag{3.16}$$

3.2.3 The management problem

Integrating all of the model components described above, the manager's decision problem is to choose the intensity of the two restoration actions in each period to maximize the expected present value of long-run net benefits conditional on the dynamics:

$$\begin{aligned} \max_{A_{X,t}, A_{P,t}} \sum_{t=0}^{\infty} \beta^t \cdot E_Z \left[\pi(H_t, A_{P,t}, A_{X,t})\right] \\ s.t.\ X_{t+1} = x(X_t, P_t, R_t, A_{X,t}|Z_t) \\ P_{t+1} = p(X_t, P_t, R_t, A_{P,t}) \\ R_{t+1} = r(X_t, P_t, R_t) \\ H_t = h(X_t, R_t) \\ Z_t \sim \text{Bernoulli}(\alpha_i), \end{aligned} \tag{3.17}$$

where β is the social discount factor. Given that the MHW frequency regime $i \in [low, medium, high]$ is known, we solve for the optimal management policy over an infinite time horizon, $t \in [0, \infty)$.

The dynamic restoration problem can be expressed as an infinite-horizon Bellman equation:

$$V(X_t, P_t, R_t) = \max_{A_{X,t}, A_{P,t}} \pi(H_t, A_{X,t}, A_{P,t}) + \beta E_Z[V(X_{t+1}, P_{t+1}, R_{t+1})], \tag{3.18}$$

subject to the same constraints as the problem in Equation (3.17) and where the expectation of the value

function on the right-hand side is given by

$$E_Z[V(X_{t+1}, P_{t+1}, R_{t+1})] = \sum_{Z_t \in \{0,1\}} Pr(Z_t) \cdot V(x(X_t, P_t, R_t, A_{X,t}|Z_t), p(X_t, P_t, R_t, A_{P,t}), r(X_t, P_t, R_t)) \tag{3.19}$$

$$\begin{aligned} = \quad & \alpha_i \cdot V\left(x(X_t, P_t, R_t, A_{X,t}|Z_t = 1), p(X_t, P_t, R_t, A_{P,t}), r(X_t, P_t, R_t)\right) + \\ & (1 - \alpha_i) \cdot V\left(x(X_t, P_t, R_t, A_{X,t}|Z_t = 0), p(X_t, P_t, R_t, A_{P,t}), r(X_t, P_t, R_t))\right). \end{aligned}$$

In reality, it is likely that the true MHW regime is unknown at first. We assess the sensitivity of our results to this regime uncertainty in our results section, considering alternative belief specifications over the prevailing MHW-frequency level.

3.2.4 Parameterization

We identified plausible parameter values for the model using a combination of peer-reviewed literature, government reports, and personal communications with experts. We also used Approximate Bayesian Computation (ABC) to identify biological model parameters that are consistent with survey data from kelp forest system surveys. Specifically, we used 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 m^2, which we mirror with the size of our representative patch model. Our implementation of ABC to identify biological parameters aligned with the observed species dynamics is 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. Parameter definitions, values and sources are summarized in Table 3.1.

3.3 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

	Description	Value	Source
γ_P	Per-unit active grazing intensity of P on X	0.00036	ABC
γ_R	Per-unit active grazing intensity of R on X	0.00068	ABC
σ_1^P	Inversely determines kelp density at maximum P active grazing on X	4.9909e-05	ABC
σ_2^P	Magnitude of active grazing of P on X	0.00484	ABC
σ_1^R	Inversely determines kelp density at maximum R active grazing on X	5.3577e-05	ABC
σ_2^R	Magnitude of active grazing of R on X	0.00495	ABC
μ_P	Per-unit passive grazing intensity of P on X spores	0.00034	ABC
μ_R	Per-unit passive grazing intensity of R on X spores	0.00081	ABC
ϕ	Beverton–Holt density-dependence for recruitment	0.00044	ABC
ϵ_1^P	Conversion factor for reproduction of P via grazing	1.5901	ABC
ϵ_1^R	Conversion factor for reproduction of R via grazing	2.7056	ABC
ϵ_2^P	Conversion × drift kelp reproduction-loss ratio (for P)	0.8748	ABC
ϵ_2^R	Conversion × drift kelp reproduction-loss ratio (for R)	0.5041	ABC
υ	Per-unit spore production of X	1.2258	ABC
τ_P	Per-unit passive grazing intensity of P on drift X	0.00322	ABC
τ_R	Per-unit passive grazing intensity of R on drift X	0.00264	ABC
ξ	Fractional mortality of X during MHW	0.4	Personal communication
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 communication
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	53 urchins/patch	Inferred from Ward et al. (2022)
R_{exo}	Exogenous recolonization of R	25 urchins/patch	Inferred using the same procedure as 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 unit of R	\$3.40/urchin	CDFW (2020)
θ_X	Unit cost of outplanting effort for X (\$/hr)	\$42.50/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.; Fishman (2022)
κ_R	Revenue multiplier from R fishery to community	3.1	Jacobsen et al. (2014)
κ_X	Economic value per unit of X	\$1.63	Assumed (equated to urchin fishing value)
β	Discount factor	0.97	Assumed
α	MHW occurrence probability for each regime $i \in \{low, medium, high\}$	{0.04,0.16,0.27}	McPherson et al. (2021); NOAA (2025)

Note: ABC = Approximate Bayesian Computation; MHW = marine heatwave. With a slight abuse of notation, here we use our variables for the population states to also refer to the species generally.

Table 3.1: Parameter definitions, values, and sources

state variables, it is not possible to visualize one of these functions in a single plot. Instead, we show three heatmap-style plots, 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 our upper bound for each variable. This degraded case is given by kelp at 25%, purple urchins at 75%, and red urchins at 25% of their maximum population level. Defining script variables as indicating percentage of the maximum population, this degraded case is given by $(\mathcal{X} = 25, \mathcal{P} = 75, \mathcal{R} = 25)$. The second alternative we consider is a restored case: $(\mathcal{X} = 75, \mathcal{P} = 25, \mathcal{R} = 75)$. Finally, we consider the midpoint between degraded and restored cases: $(\mathcal{X} = 50, \mathcal{P} = 50, \mathcal{R} = 50)$.

In Table 3.2 we show the model components that allow for alternative specifications. We begin in the next section by presenting results for our "baseline scenario" as indicated by the underlined specification in the table. After establishing this baseline, we present results for alternative scenarios, focusing on results that differ from the baseline.

Component	**Alternatives**		
MHW regime	low	medium	high
MHW regime knowledge	uncertain	certain	
Red urchin harvest	not allowed	allowed	
External kelp recruitment	available	not available	
Restoration actions used	A_X only	both	A_P only
Policy form	static proportional heuristic	optimal dynamic feedback	
Latent states in plotting*	degraded	midpoint	restored

Note: Our "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.

Table 3.2: Model specification alternatives

3.3.1 Baseline scenario

In Figure 3.1 we show the value function for our baseline scenario, which reflects an infinite time horizon perspective, a medium MHW intensity regime that is known with certainty, and red urchin harvesting allowed. The solution is shown at the midpoint case for latent states. As expected, the value function is increasing in kelp population and red urchin population and decreasing with purple urchin population. Red urchins still provide a net positive long-term benefit, despite the tradeoff they pose, i.e., fishery benefits versus

loss of direct kelp value from active grazing on kelp. The marginal net benefit of additional red urchins is stronger given high kelp population (first subfigure), indicating that, as expected, healthy kelp facilitates the contribution of red urchins to welfare. Conversely, the middle subplot shows that the marginal net benefit of red urchins diminishes as purple urchin population rises, though it remains positive. This pattern confirms that under the heavy active grazing pressure of high purple urchin populations, the desirability of additional red urchins is diminished.

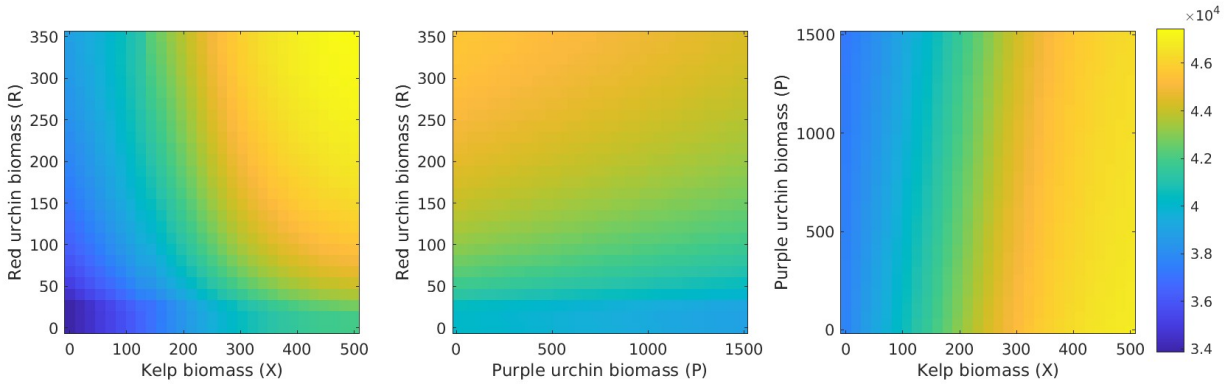
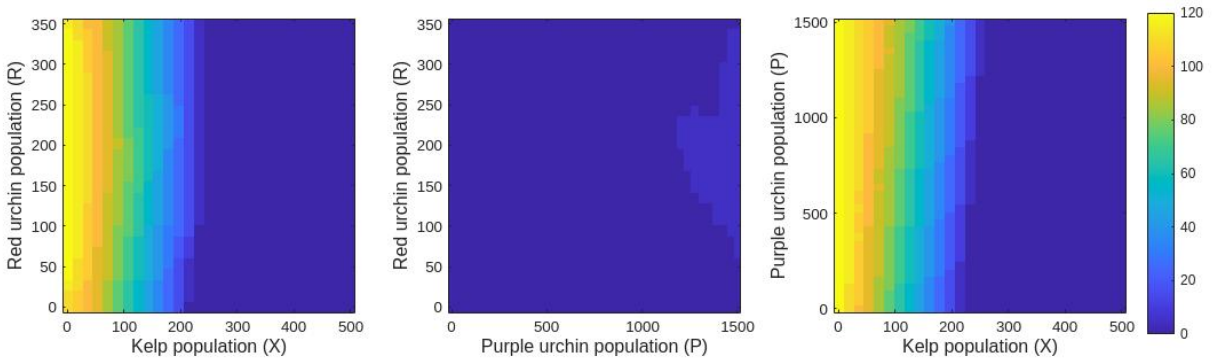


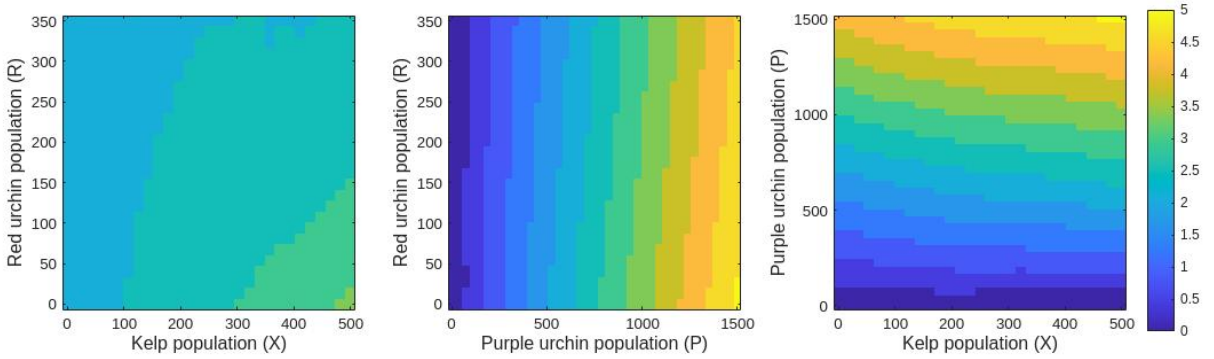
Figure 3.1: Value function for the baseline scenario specified in Table 3.2, shown as a heatmap for each pair of state variables on the vertical and horizontal axes, with latent states at their midpoint.

Figure 3.2 shows the policy functions for the intensities of kelp outplanting (A_X, 3.2a) and the removal of purple urchins (A_P, 3.2b). Optimal outplanting varies only modestly with both urchin populations: A_X increases with purple urchin abundance (P) and increases then decreases (again, modestly) with the red urchin population (R). A_X varies only weakly with P because this purple urchin population is simultaneously being controlled by removals, A_P (as evidenced in results shown further below in the scenario where A_P is restricted to zero). The dominant driver of changes in the efficient level of outplanting is intuitive: A_X increases strongly as kelp population (X) declines. When kelp is scarce, outplanting yields high marginal benefits through multiple channels: direct value of additional kelp, increasing the stock that can naturally reproduce, and driving the kelp stock to levels with desirable urchin outcomes where urchins subsist on drift kelp (rather than kelp stipes) and where red urchin gonad quality facilitates fishing. Once kelp population recovers to this zone—just above $X = 200$, safely above $X^P_{\text{max-graze}}$—natural kelp recruitment sustains the system and efficient outplanting ceases.

In contrast to outplanting, the optimal purple urchin removal intensity (A_P), shown in Figure 3.2b, is generally non-zero except when purple urchin population (P) is very low. As expected, the dominant



(a) Optimal kelp outplanting intensity, $A_X(X, R, P)$



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

Figure 3.2: Optimal restoration levels (labor hours) for the two management options, kelp outplanting (top row) and purple urchin removal (bottom row), for the baseline scenario, as specified in Table 3.2.

driver of optimal removal levels is the population of purple urchins, with effort increasing consistently as P increases. Efficient removal is only weakly increasing as red urchins increase. Since both red and purple urchins may feed directly on adult kelp stipes, we might expect that a higher level of R would create a strong incentive to reduce the overall direct grazing pressure by lowering P. However, red urchins are only a secondary active consumer of kelp stipes (see Equation (3.6)). Furthermore, red urchins compete with purple urchins for food—in particular, red urchins are a primary consumer of drift kelp, before purple urchins (see Equation (3.8)). Thus, a larger red urchin population generates a check on purple urchin reproduction. This weakens, but does not eliminate, the incentive to expand removal of purple urchins as grazing pressure grows with red urchins.

While we might expect the desirability of removals to increase as the kelp population deteriorates, instead, we find the opposite. As shown in the first and final subplots of Figure 3.2b, optimal removals actually fall weakly as kelp contracts. This is particularly surprising since, as kelp declines from abundant levels, the rate of active grazing per urchin increases (following the Type IV functional response in Equation (3.2)).

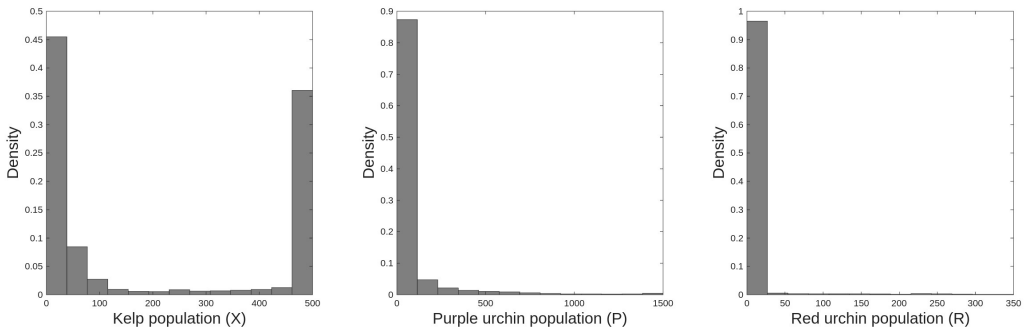
However, at the same time, the amount of kelp subject to grazing—the value at risk—is also falling. Here, we find that the latter effect dominates, albeit weakly.

Figure 3.3 presents density functions for kelp and urchin population levels after 50 years, under no restoration (3.3a) or with optimal restoration given by the policy functions in Figure 3.2 (3.3b). The starting state in year one was our midpoint state case ($\mathcal{X} = 50, \mathcal{P} = 50, \mathcal{R} = 50$), however, after the 50-period horizon applied here, the distributions are not sensitive to this initialization. Under no restoration, Figure 3.3a shows that both purple and red urchins are most likely to remain at low population levels, and only infrequently reaching high population levels. This pattern is driven primarily by the scarcity of kelp as a food source, e.g., due to occasional MHW events. The kelp population distribution is bimodal, with probability mass concentrated at low and high densities. Here, we highlight the fact that under the baseline model, in particular a medium intensity MHW regime, without continuing restoration effort, the red urchin population will not sustain a fishery.[9] While the univariate densities in Figure 3.3 are straightforward, they do not show how state levels are likely to co-vary across the three species. To show this, we present bivariate densities in heatmap form in Appendix Figures D13b (no restoration) and D14b (optimal restoration). Under no restoration, we find that when the kelp population is high, urchin populations tend to be very low.

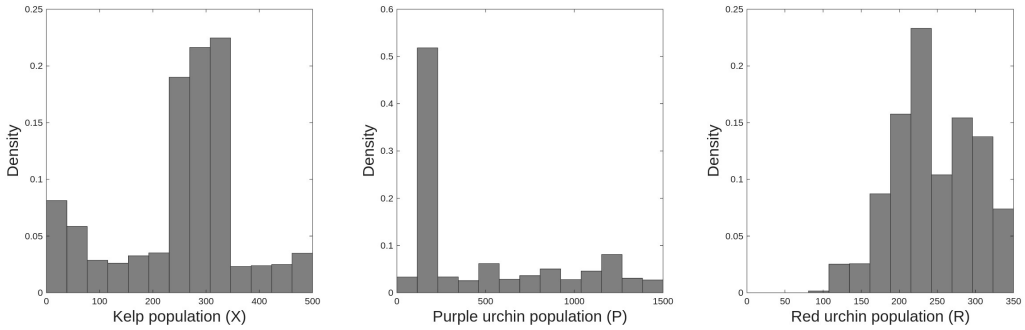
In contrast to the no restoration case, Figure 3.3b shows that under optimal restoration (outplanting and urchin removal) kelp and red urchin populations are likely to be at higher levels, while purple urchins likely remain suppressed, in this case due to targeted removal instead of low kelp levels. In this ongoing restoration case, the bivariate heatmap densities in Appendix Figure D14b show that high kelp population commonly co-occurs with low purple urchin and *high* red urchin levels.

Next, we consider the set of alternatives to our baseline scenario summarized in Table 3.2. To aid comparison, Table 3.3 provides a summary of top-line results for each alternative scenario, specifically how value and policy function results change relative to the baseline scenario above. In the main text we focus on results that depart notably from the baseline and show remaining results in the Appendix.

[9] In contrast, when the MHW regime is low and there is no restoration, the red urchin population is generally high. We explore this alternative scenario in more detail below.



(a) Restoration: none



(b) Restoration: optimal policy

Figure 3.3: Probability mass 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, conditional on the given policy, over 50 periods and integrating over the other two state variables. The starting state was our midpoint state case ($\mathcal{X} = 50, \mathcal{P} = 50, \mathcal{R} = 50$).

3.3.2 Biology and climate scenarios

Alternative scenario 1: Latent states in plotting that represent either degraded or restored cases

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) was set to its midpoint. Here, we assess whether our qualitative results are sensitive to this choice by considering results selected from alternative latent state cases. These alternative cases do not represent changes to the model itself but rather expanding the viewpoints taken of the solution, specifically, alternative two-dimensional slices of the three-dimensional value and policy functions. The first alternative case we consider 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 ($\mathcal{X} = 75, \mathcal{P} = 25, \mathcal{R} = 75$), as defined at the beginning of Section 4.3. These results are shown in Appendix Figures D6, D7 and D8. They show that numerical magnitudes of the value

Factors	Scenario		Value function		Outplanting (A_X)		Removal (A_P)	
			Shape	Level	Shape	Level	Shape	Level
Biology & climate	1. Latent states	degraded state	·	–	·	+	·	+
		restored state	·	+	·	–	·	–
	2. MHW-frequency regimes	low MHW regime	·	+	·	–	·	·
		high MHW regime	·	–	·	+	·	·
	3. Kelp spore loss & exog. recruitment		·	– (~6–8%)	·	+	·	·
Social & management	4. Heuristic policy: static proportions		N.A.–level drop dominates	– (~90%)	N.A.–level drop dominates	– (drops to 0)	N.A.–level drop dominates	– (drops to very low level)
	5. No red urchin harvest		Decr. in R (instead of incr.)	– (~50%)	Drop more pronounced at low R	–	·	·
	6. Restricted restoration	outplant (A_X) only	·	– (20–30%)	Incr. more pronounced as P rises	+	N.A.	
		removal (A_P) only	Decline sharper at low X with non-low R	– (0–100%)	N.A.		·	Minor shifts; zero near $X = 0$

Note: All entries describe changes relative to the baseline. Key: "·" = no meaningful change from baseline; + / – = increase/decrease in level; for level shifts that are either fairly uniform or stark, we indicate the relative magnitude. N.A. = not applicable.

Table 3.3: Summary of results across alternative scenarios relative to the baseline.

and policy functions can shift intuitively, e.g., when the latent state of kelp is higher (restored case), the value function shifts higher and optimal kelp outplanting is lower. But overall the qualitative results from our baseline model results in Section 3.3.1, from the vantage point of midpoint latent state case, are unchanged.

Alternative scenario 2: MHW-frequency levels and uncertainty

Relative to the baseline scenario, which assumes a medium MHW frequency, we evaluate how low- and high-MHW frequency regimes influence the results. For value, overall patterns are unchanged and results simply reflect the fact that net returns are higher under the reduced thermal stress of the low-MHW regime and, conversely, lower under the high-MHW regime (see Appendix Figure D9).

For optimal outplanting policy, overall patterns are also unchanged from the baseline model results summarized above, though there is again a level shift. As expected, when average kelp mortality is lower given a low regime, optimal kelp outplanting falls, and there is a small contraction of the kelp population range over which this effort is non-zero. Conversely, under a high regime, effort increases and there is a small expansion in the range over which some effort is prescribed (see Appendix Figure D10).

We find that the removal intensity is insensitive to the MHW frequency (see Appendix Figure D12). This

is due to the fact that the direct biological impact of MHWs is on kelp survival (not purple urchins) and the kelp outplanting policy discussed above is available to directly respond to those changes in kelp mortality.[10]

The distribution of long-run levels of the three species show interesting sensitivity to the MHW regime. First, in the *absence* of any restoration effort and with low MHW frequency—consistent with conditions in the early 2000's—all species have a substantial likelihood of avoiding low levels (see Appendix Figure D15). But when the regime is high—reflecting conditions from 2010-2020—and there is still no restoration, both urchin populations are consistently very low and kelp is strongly bimodal—either very low (due to high MHW mortality) or very high (given minimal urchin active grazing). In contrast to this sensitivity to MHW regime in the absence of management, the introduction of optimal restoration blunts population sensitivity to climate extremes (see Appendix Figure D16). But sensitivity is not fully eliminated: relative to the baseline model (Figure 3.3b), when MHW frequency is low, all stocks have less spread and are concentrated around (desirable) modal levels. Conversely, when the regime is high, the distribution of stock levels shows more spread, particularly in undesirable directions.

We also examine how optimal policies change when restoration managers are uncertain about the true MHW regime. We define the ϕ as the vector of probability weights placed on each of the possible true MHW-frequency levels, $i \in \{low, medium, high\}$. We consider three belief specifications: (1) uniform, $\phi = \{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\}$; (2) low-frequency skewed, $\phi = \{\frac{2}{3}, \frac{1}{6}, \frac{1}{6}\}$; and (3) high-frequency skewed, $\phi = \{\frac{1}{6}, \frac{1}{6}, \frac{2}{3}\}$. Under each belief specification, the expected MHW occurrence probability is the weighted sum, $\sum_i \phi_i \alpha_i$. Since the uniform case generates an expected MHW occurrence equal to $i = medium$, the results are equivalent to the baseline scenario. Across the two skewed-belief cases, the purple urchin removal policy remains nearly unchanged (as found for the pure low- and high-frequency regimes earlier in this section). However, the outplanting policy shifts towards that of the low- or high-frequency regime to which beliefs are skewed (Appendix Figure D11). Thus, overall we find that purple urchin removals are insensitive to MHW-frequency levels or uncertainty. In contrast, outplanting differs at the two frequency extremes and scales between as beliefs shift—from skewed low to high–between the two extremes.

[10]There is one result that departs substantially from the baseline policy, which we do not emphasize in the main text since it emerges (1) as an artifact of the purple urchin population upper bound and (2) in a part of the state-space almost never visited. Under the low–MHW regime, a small region with no purple urchin removal appears at high kelp (X) and high purple urchin (P) levels (see Appendix Figure D12a). In that corner of the state space with high X and P, favorable kelp recruitment under a low-MHW regime would allow P to grow beyond the upper bound for purple urchins; however, we force such outcomes back to the upper bound. As such, the marginal benefit of removal is zero and the optimal policy is $A_P = 0$. The density of state variable outcomes after 50 periods shows that the likelihood of visiting the state space zone of this artifact is exceedingly small (see Appendix Figure D14a).

Alternative scenario 3: Kelp spore loss and exogenous recruitment

In our baseline model, kelp dynamics are closed: no spores disperse to, or enter from, surrounding waters. This is also consistent with a setting in which surrounding patches experience the same conditions as our focal patch and mixing in the water column is uniform, such that any outflow of spores is equal to the inflow. Here we relax this assumption to explore the effect of external spore arrival, dispersing from remote kelp refugia a few kilometers away (2-10 km). Such patches of persistent kelp have been observed, for example, in California near river mouths where lower salinity is tolerated better by kelp than by urchins (Ricart et al., 2025). To incorporate this feature, we modify kelp dynamics to add external spore recruitment from a site that produces a constant level of spores, of which the share ζ settle in our focal patch. Symmetrically, spore production within our focal patch disperses outside of the patch at the same rate, ζ (see Appendix A.2 for details). We take this symmetric approach, with both entry *and* exit of spores, so that results mainly reflect a shift in the nature of recruitment and not, additionally, an overall increase in spore production. We select levels for the external recruitment share ζ, and the recruitment strength of the external site by re-calibrating this extended model to the data (see Appendix C.1).

Under this kelp spore loss and external recruitment, we find no substantial difference in optimal urchin removals. In contrast, optimal outplanting intensity increases overall, as shown in Appendix Figure D18. To understand why it is efficient to increase kelp outplanting when external recruitment of kelp is available, note this change to the model results in a net increase in natural recruitment of kelp when the in-patch adult kelp level is low but also a net *decrease* when the in-patch adult kelp level is high. The threshold between these two cases occurs when kelp is around 20 to 40, above which spore outflow exceeds inflow. Thus, over most of the range for X, the calibrated model drives net spore loss. As a result, the level of optimal outplanting increases.

Relative to the baseline, the value function (Appendix Figure D17) shifts downward, with expected welfare levels about 6-8% lower. We find that the distribution of likely future state levels shifts noticeably under kelp spore loss and external recruitment, as shown in the long-run probability density functions in Figure 3.4 for both the baseline and this alternative scenario. While the modes and overall shapes for each species are broadly similar, under kelp spore loss and external recruitment we find that kelp and purple urchins exhibit modestly greater spread, both reaching higher levels with greater likelihood. This is intuitive, since exogenous kelp recruitment buffers against lower in-patch recruitment years and additional

kelp supports additional purple urchins. Red urchin levels shift slightly lower. One potential channel is food competition with purple urchins through the direct consumption pathway. Although kelp is higher on average in this scenario, we find that red urchins' direct kelp consumption falls relative to the baseline, while purple urchins' direct kelp consumption rises relative to the baseline. This is consistent with higher purple urchin abundance reducing red urchins' realized intake despite greater standing kelp population, which slows red urchin reproduction and increases the likelihood of lower population states relative to the baseline.

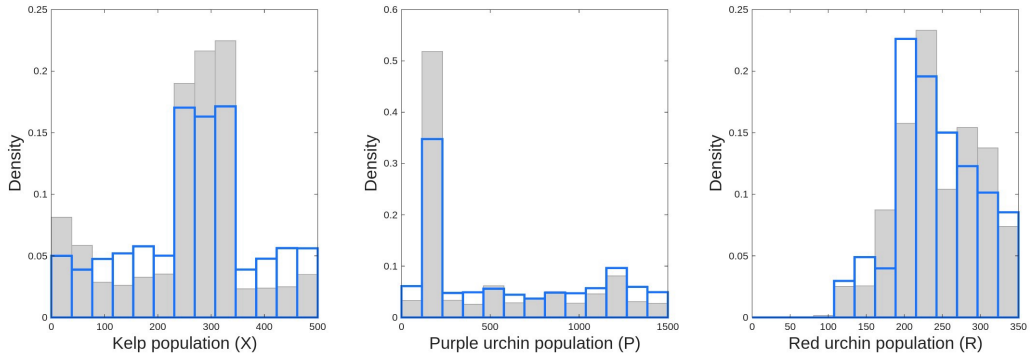


Figure 3.4: Probability density functions for state variable levels after 50 years under the baseline scenario (gray) and under the addition of kelp spore loss and exogenous recruitment (blue), given optimal restoration policy.

3.3.3 Social and management scenarios

Alternative scenario 4: Restoration set by a static proportional heuristic policy

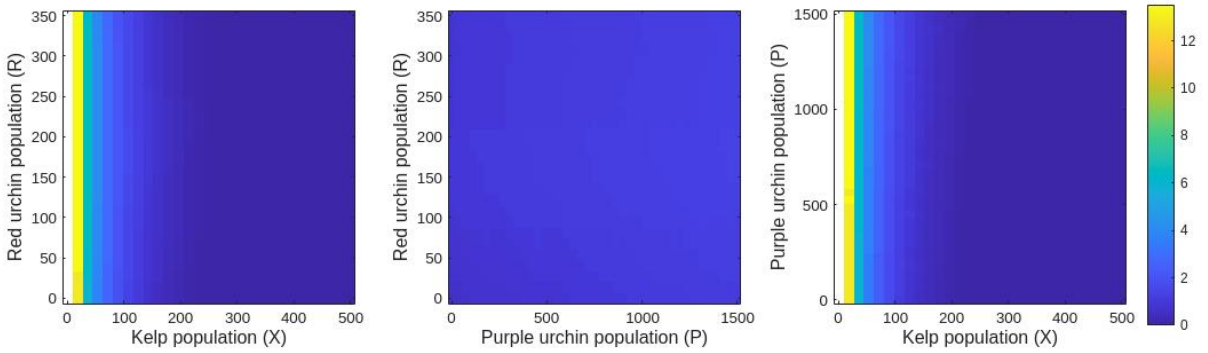
Implementing state-dependent, feedback-controlled optimal restoration, as given by our policy function results, can be computationally challenging. A heuristic, rule of thumb approach, such as holding relative restoration intensities fixed over time, offers a simpler alternative, albeit with efficiency losses. For example, Arroyo-Esquivel et al. (2023) evaluate biological outcomes under fixed proportional restoration policies, i.e., removing a fixed fraction of purple urchins and outplanting a fixed fraction of kelp each period.[11] Here we consider a scenario in which the manager adopts such a static proportional policy.

We specify static proportional restoration intensities (b_X, b_P), each on the unit interval, for kelp outplanting and purple-urchin removal, respectively. Each period, the manager outplants $b_X X_t$ kelp stipes and removes $b_P P_t$ purple urchins. To identify the best-performing proportions (b_X, b_P), we consider a two-dimensional grid spanning both parameters. Assuming the starting state is the degraded state case,

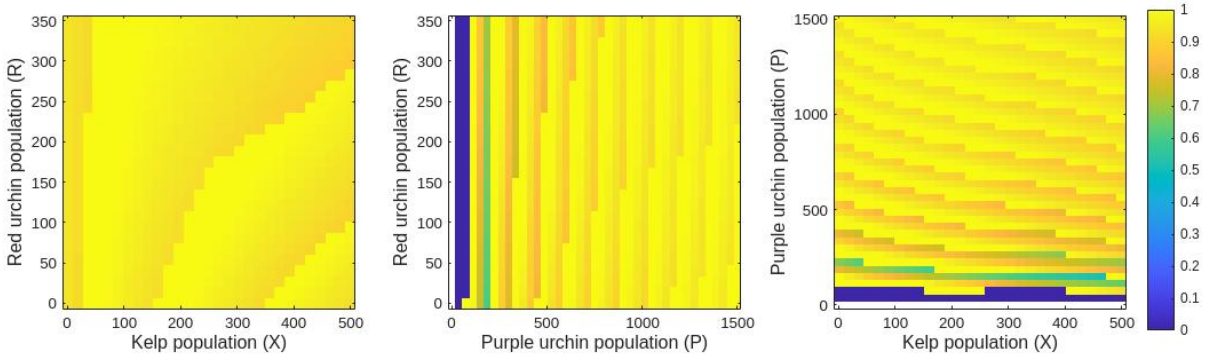
[11] Arroyo-Esquivel et al. (2023) focus on biological outcomes and do not include economic components.

consistent with current Northern California conditions, for each pair (b_X, b_P) we consider the value function result (infinite-horizon discounted expected net benefits) as shown in Appendix Figure D20.

The best-performing static proportional policy is $(b_X, b_P) = (0, 0.04)$: outplanting falls to zero and the purple urchin removal rate is 4% of the population, annually. This results in much lower long-run net benefits from the system. The value function level is roughly one-tenth the level obtained under the baseline optimal policy.[12]



(a) Optimal kelp outplanting proportion (converted from $A_X(X, R, P)$)



(b) Optimal purple urchin removal proportion (converted from $A_P(X, R, P)$)

Figure 3.5: Optimal restoration proportions for the two management options: kelp outplanting (top row) and purple urchin removal (bottom row). Values are converted from restoration effort hours into proportions relative to the corresponding population levels in the baseline scenario (except latent states in plotting are degraded). Regions with $X = 0$ or $P = 0$ are excluded from the plot (shown in white), because the ratio is undefined at zero population.

To further explore why static proportions perform so poorly in this case, we transform our policy function

[12] Because heuristic restoration policies are typically proposed for use only when the system is in distress—rather than in all states as modeled above—we considered an alternative form of the static proportional policy that operates only when the kelp population is low, specifically, below $X = 200$, which is the level below which we find optimal outplanting scales up in our baseline model policy function (Fig. 3.2). This makes little difference to the best-performing static rule: outplanting remains zero and urchin removal intensity shifts only slightly higher, from 0.04 to 0.05. Adding this low-kelp trigger does not improve the static policy long-term net benefits, which are actually even lower than those without the trigger.

from levels (see Figure 3.2) to proportions as shown in Figure 3.5. For outplanting (3.5a) the pattern is as expected: extremely low kelp population yields very high optimal restoration proportions, which decline to 0% as kelp increases. This stark gradient to the optimal proportion outplanted cannot be well-approximated by a static proportion, and thus best set to zero. For removals (3.5b) another distinctive pattern emerges: the optimal proportion generally corresponds to removing the vast majority of purple urchins, except in regions where P is low.

For proportional purple urchin removals, the large gap between the static and optimal dynamic policies arises from the interaction between the two restoration actions: intensive purple urchin removal is ideal only when paired with selective kelp outplanting. Under the static approach, when $b_X = 0$, aggressive removal (e.g., $b_P \in [0.85, 1]$) reduces active grazing pressure yet leaves the system mired in states with very low kelp and red urchins; kelp recovers only slowly through natural dynamics, and returns to aggressive removals are low. The resulting best static policy features a low removal rate ($b_P \simeq 0.04$). By contrast, the optimal dynamic policy calls for substantial outplanting (i.e., high b_X) under kelp scarcity, lifting the system out of the low-kelp trap and driving higher returns to consistently aggressive removals.

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

While our baseline model includes both direct benefits from kelp as well as profits from the red urchin commercial fishery, it is possible that even after restoration, this fishery might not return after prolonged inactivity. This might occur, for example, due to loss of specialized skills from diver exit, and/or loss of specialized diving and boat equipment. In contrast to the baseline model, where red urchin harvest occurs whenever gonad quality is high (see Equation (3.11)) here we consider an alternative scenario in which there is no red urchin harvest in any conditions.

As shown in Figure 3.6, we find that optimal outplanting intensity shifts down, as expected. But the drop is especially pronounced for low red urchin populations. This indicates that fishery profits incentivize additional outplanting investment when red urchin populations are weak, in order to boost future harvest. We also find that optimal purple urchin removal is insensitive to the exclusion of the fishery, as shown in Appendix Figure D22. This indicates that the level and pattern of the removal policy are essentially driven by direct kelp benefits.

The two main effects in the value function are both as expected (see Appendix Figure D21). First, since in the baseline model the direct kelp benefits and the red urchin fishery benefits are of similar magnitude,

when the latter are omitted the value falls by approximately half. Second, the value shifts from increasing in the red urchin population to decreasing since there are no harvest benefits but active grazing pressure remains.

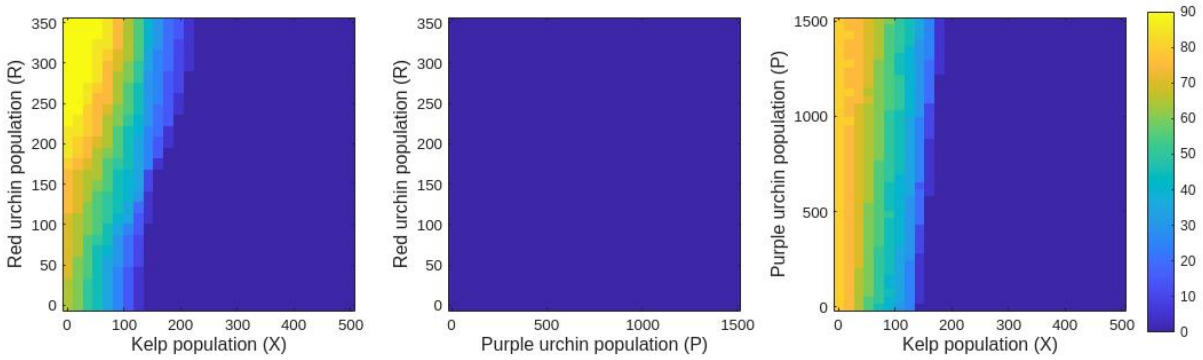
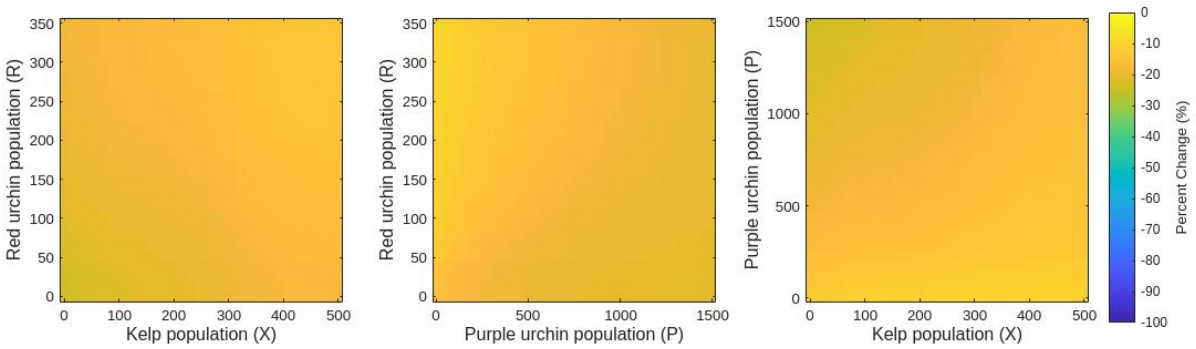


Figure 3.6: Optimal kelp outplanting level given baseline model except with no return of red urchin harvest.

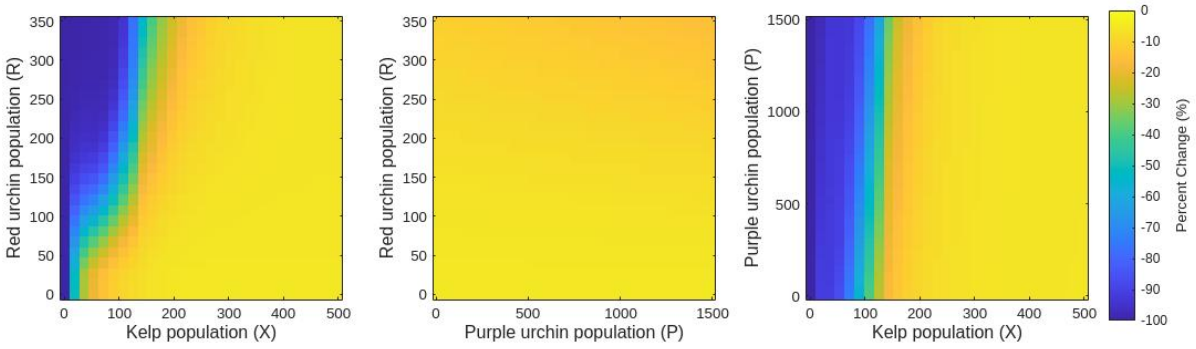
Alternative scenario 6: Restricted restoration

Our baseline model solves for the optimal joint deployment of both restoration activities (outplanting and purple urchin removal) and the results demonstrated that simultaneous use of some level of both actions is ideal over much of the state space. However, given the practical challenges of coordinating multiple management actions, in practice, such restoration may focus on a single tool. Here we investigate how the sole use of a single action changes its ideal deployment and whether such a restriction generates minor or major losses in long-run net benefits.

We show the percent change in value functions from restricting restoration to a single action in Figure 3.7. Specifically we plot $100\% \times (V^{\text{rest}} - V^{\text{base}})/V^{\text{base}}$, where V^{rest} is the value function under restricted restoration and V^{base} is from the baseline (unrestricted) scenario. The pattern of losses varies greatly between the two alternative restrictions. Using outplanting only (3.7a) results in fairly consistent expected losses of around 20-30% in this midpoint case, easing slightly when purple urchins are low and red urchins are high. In contrast, using removals only (3.7b) can lead to either negligible or substantial loss. Stark losses from eschewing outplanting intuitively occur when kelp is low, but ease when red urchins are also low, since forgone fishery benefits would be lower in this case. Overall, we find that such a restriction can reduce the long-run net benefits anywhere from negligibly to substantially, depending on the restriction and population state levels.



(a) Implementing only kelp outplanting, A_X (urchin removal is restricted, $A_P = 0$)



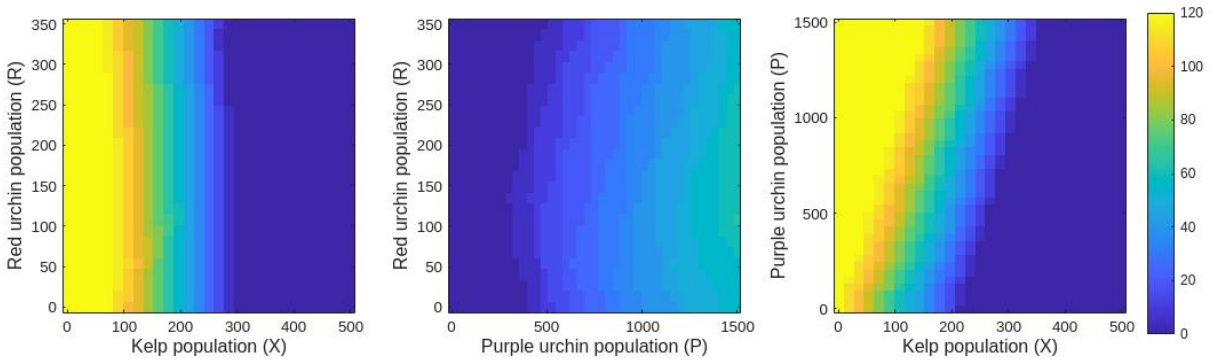
(b) Implementing only purple urchin removal, A_P (kelp outplanting is restricted, $A_X = 0$)

Figure 3.7: Percent change in the value function under restricted restoration relative to the baseline scenario: outplanting only (top, $A_P = 0$) and purple urchin removal only (bottom, $A_X = 0$).

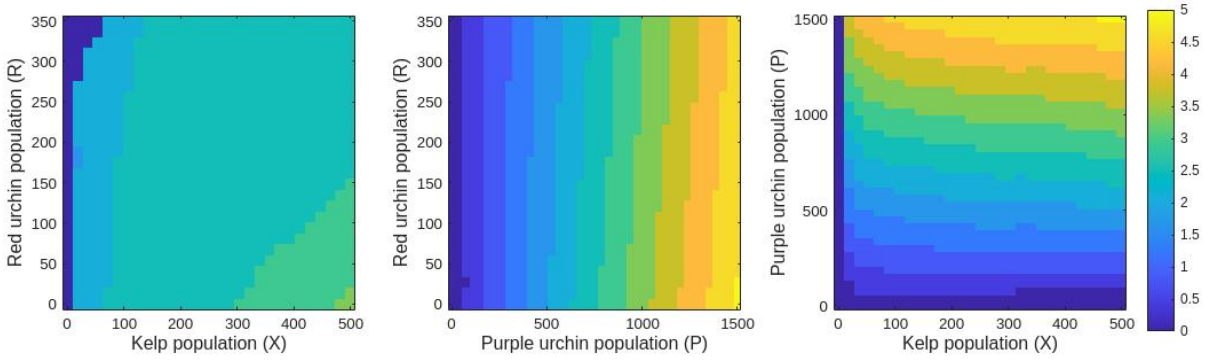
We show the policy functions under restricted restoration scenarios in Figure 3.8: kelp outplanting only (top row) and purple urchin removal only (bottom row). In Figure 3.8a, the kelp outplanting intensity shifts noticeably across the state space compared to the baseline scenario (Figure 3.2), showing substantial expansion of restoration effort application across states and intensification of restoration efforts as well. By contrast, when only purple urchin removal is implemented (3.8b), the changes in policy intensity are relatively minor.[13] Overall, we find that outplanting adjustments can partially compensate for the non-use of removals and thus partially moderate restriction losses; this is not true in the alternative case in which removals alone are used.

Finally, we assess whether the particularly poor performance of using removals alone might be ameliorated if there is also exogenous kelp spore loss and recruitment as modeled in Section 3.3.2. We do not find meaningful differences in optimal urchin removal policy under external recruitment (Appendix

[13]One apparent difference from the baseline scenario emerges when kelp population is near zero where there is a small no-restoration zone. Once the system arrives at a near-zero kelp state, kelp recovery without outplanting becomes very difficult and investment in urchin removal is much less beneficial.



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



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

Figure 3.8: Optimal restoration levels given the baseline model except with restricted restoration: outplanting only (top, $A_P = 0$) and removal only (bottom, $A_X = 0$).

Figure D24b).[14] Considering the long-run net benefits, we do find that at very low kelp levels external recruitment buffers against kelp collapsing towards zero and thus against the very low values experienced there. However, away from very low kelp states, this kelp spore connectivity does not compensate well for the absence of outplanting since spore losses become salient: value levels are lower across most of the state space compared to the case of no spore connectivity discussed just above (see Appendix Figure D23).

3.4 Discussion

In this study we develop a dynamic bioeconomic model of kelp-urchin system that is the first to identify welfare-maximizing restoration actions that flexibly depend on the current state of the system. The decision framework identifies the optimal mix of two restoration options and integrates ecological feedbacks, fishery

[14] We do observe the emergence of a small no-restoration zone, similar to the artifact in the urchin removal policy under low MHW (Appendix Figure D12a), which arose from the upper bound on purple urchin population. Here, we confirm that the artifact is generated by the upper bound on red urchin population. Under restoration, these corner states (low X, high R simultaneously) are almost never visited, so this does not affect our results.

dynamics, restoration costs, socio-economic benefits, and marine heatwave (MHW) uncertainty.

We establish results for a baseline model, then assess when and how alternative modeling assumptions generate results that substantially deviate. In our baseline scenario, we find stark differences in the dynamic use of the two restoration tools. Optimal kelp outplanting is critical when the kelp population is low but declines steadily to zero as kelp increases to medium levels. In contrast, purple urchin removal is non-zero and quite aggressive across nearly all state conditions (except when the purple urchin population is very low). While most baseline insights are intuitive, some surprising results emerge: though we might expect ideal restoration intensity should generally increase as conditions worsen, we find that optimal removals weakly decrease as kelp contracts.

We summarize how alternative scenario results follow or deviate from baseline model results in Table 3.3. Here, we highlight four alternative scenario insights that stand out. First, we find that ideal use of one restoration action is quite sensitive to the specification of MHW frequency, while the other is not. Purple urchin removal is insensitive to both MHW-frequency levels and uncertainty about the prevailing MHW regime. Outplanting, in contrast, differs across the low and high frequency extremes considered, and scales between them as MHW beliefs under regime uncertainty shift from low-frequency-skewed to high-frequency-skewed. This asymmetry arises intuitively because MHWs act directly on kelp mortality, so adjustments fall primarily on the kelp-augmenting restoration tool. We also find that, restoration blunts population sensitivity to climate extremes: without restoration, kelp outcomes are bimodal (very low or very high) and urchin populations in both species remain chronically low (thus no harvest as well), whereas applying optimal restoration results in long-run outcomes concentrated around desirable state levels of high kelp, low purple urchins, and viable red urchins for harvest).

Second, even if red urchin fishery recovery is not feasible—e.g., due to loss of specialized knowledge and equipment—other direct benefits of kelp can still motivate substantial investment in restoration. Relative to the baseline model, in which fishery benefits and direct kelp benefits are comparable, exclusion of the fishery leads to less use of kelp outplanting, but purple urchin removals are essentially unchanged. This suggests that uncertainty over the return of the commercial fishery is not a core concern for purple urchin removal strategy, though it motivates modulating outplanting.

Third, replacing the state-contingent feedback policy with a simpler heuristic in the form of a static proportional decision rule leads to drastically different restoration use and major loss of value. This large gap arises from the interaction between the two restoration actions. The optimal feedback policy involves

consistent intensive purple urchin removal but much more variable outplanting effort that drops from high to zero as the kelp stock increases. The former (removals) can be reasonably approximated by the static heuristic but the latter (outplanting) cannot. The result is that the best-performing static proportional rule involves no outplanting. This absence of outplanting drastically changes the returns to urchin removal, which are consequently set to a very low rate under the static heuristic. Thus, losses from the poor fit of a heuristic to one action (outplanting) compound as this interacts with the action for which the heuristic would otherwise be a reasonable fit. Of course, one could choose an alternative to the static proportional heuristic for kelp—e.g., linearly declining kelp outplanting as kelp increases—but the insight that such an alternative is a reasonable candidate is available to us because the optimal feedback policy has already been solved for.

Fourth, restricting restoration to a single tool can reduce welfare negligibly or substantially, depending on the restriction and population levels. Overall, outplanting adjustments can partially compensate for the non-use of removals, moderating the welfare loss from constraining tool choice. This is not true in the alternative case in which removals alone are used. Relying on removal alone performs very poorly when kelp is sufficiently low to prompt direct grazing from urchins. Here, eschewing outplanting can lead to a prolonged recovery path with sluggish kelp rebuilding and delayed red urchin harvest as the red urchin population rebounds slowly at low-kelp states. If grazer control is the sole, long-run tool of choice, a temporary outplanting pulse may still be necessary to boost kelp out of the direct-grazing zone; but once that is achieved, using removal alone performs well. Overall, the importance of jointly deploying both tools is context-dependent, more critical when considering using removal alone and when the system is currently degraded.

While our analysis considers several alternative scenarios as discussed above, several limitations remain. In the realm of spatial spillovers, we consider only two different specifications for kelp spore exchange and ignore urchin movement across patches. Different forms of spatial connectivity are likely to affect recovery dynamics and optimal restoration. Key parameter uncertainty remains, for example, the MHW-induced kelp mortality parameter is based on laboratory estimates not yet validated in the field. While we capture complex feedbacks between three key species, additional species may play a crucial role, e.g., urchin predators such as sunflower sea star (*Pycnopodia helianthoides*). Future work incorporating explicit spatial patch interactions, additional key species, and updated parameter estimates capturing evolving scientific understanding would strengthen insights into efficient kelp forest restoration strategies.

Appendices

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 (3.7) and (3.8) in the main text. In a continuous-time framework, the evolution of drift kelp, denoted by D_t, follows the law of motion:

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

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:

$$\bar{D_t} = \frac{o_1 \cdot X_t}{o_2 \cdot P_t + o_3 \cdot R_t + \Delta}. \tag{.2}$$

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:

$$\bar{D_t} = \frac{o_1 \cdot X_t}{o_2 \cdot P_t + o_3 \cdot R_t + \Delta} \approx \frac{o_1 \cdot X_t}{\Delta}. \tag{.3}$$

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:

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

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:

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

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

$$p_D(X_t, P_t, R_t) = \epsilon_3 \cdot X_t \cdot \exp(-\tau_R \cdot R_t) \cdot (1 - \exp(-\tau_P \cdot P_t)), \tag{.6}$$

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.

Under the Type IV functional response, combined with adult kelp and drift kelp grazing dynamics and our parameterized biological processes, we find that urchin grazing behavior reproduces patterns consistent with real-world observations. In particular, Figure A1 shows that, holding P and R fixed, the grazing impact of both urchin species decreases as kelp population increases. Furthermore, purple urchins generally exert the dominant grazing pressure.

Furthermore, for red urchins, Figure A2 illustrates the balance between direct adult kelp grazing and drift kelp grazing across different kelp population levels. As expected, red urchin grazing on drift kelp becomes increasingly dominant relative to direct adult kelp grazing as the kelp population rises.

A.2 Kelp dynamics with exogenous spore recruitment

We modify kelp dynamics to allow spores to exit and enter our focal representative patch. A share ζ of spores produced within the patch exit. In addition, there is an external source of spore production at rate η_E in a remote patch, of which the share ζ settle in our focal patch. The modified state-transition equation for

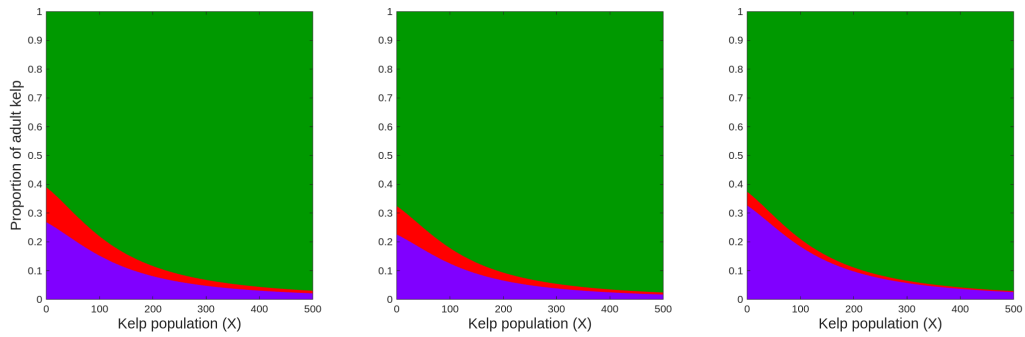


Figure A1: Share of surviving adult kelp (green), grazing share of adult kelp by purple urchins (purple), and by red urchins (red), across kelp population levels under a Type IV functional response. Each subfigure corresponds to a different fixed urchin population level: (i) the mean population density observed in the long-run analysis without restoration, (ii) the midpoint population level, and (iii) the degraded state.

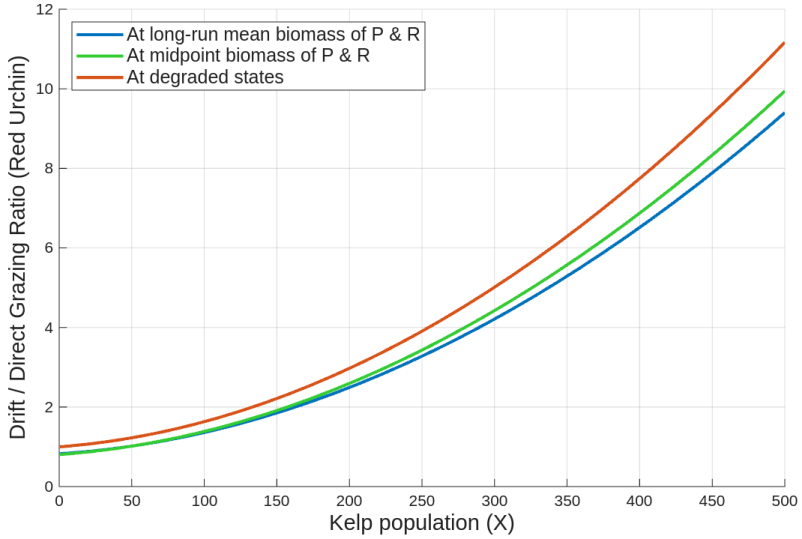


Figure A2: Relative contributions of direct adult kelp grazing and drift kelp grazing by red urchins across kelp population levels. Each line corresponds to a different fixed purple urchin population level.

kelp is, where $k_t \equiv k(X_t, P_t, R_t)$:

$$X_{t+1} = \left[\left((1-\zeta)\cdot \upsilon \cdot k_t + \zeta \cdot \eta_E\right) \cdot \frac{\exp\left(-\mu_P \cdot g_P(X_t, P_t) - \mu_R \cdot g_R(X_t, R_t)\right)}{1 + \phi \cdot k_t} \; + \; k_t \cdot \delta_X \right] \cdot \left(1 - \xi \cdot Z_t\right) \; + \; \eta_X \cdot A_{X,t}.$$

For the newly introduced parameters ζ and η_E, and all other biological parameters of the model, detailed parameter estimates are provided in Appendix C.1.

A.3 Fixed proportional policy: model modification and evaluation

We outline how the baseline model is modified to accommodate a fixed, state–independent proportional restoration policy for kelp outplanting and purple urchin removal, and how the associated value function is evaluated.

Following Arroyo-Esquivel et al. (2023), restoration is applied proportionally to current populations. In each period t, a fraction $b_X \in [0, 1]$ of the current kelp stock X_t is added via outplanting, and a fraction $b_P \in [0, 1]$ of the current purple urchin stock P_t is removed. With this proportional and fixed policy, we convert these proportions into implied hours using the efficacy parameters (η_X, η_P).

Once the proportional restoration intensities are chosen, the implied restoration hours follow from efficacy. For purple urchin removal, removing a fraction b_P of the current stock P_t requires hours $A_{P,t}$ that satisfy $\eta_P\, A_{P,t} = b_P\, P_t$. Applying the same logic to kelp outplanting yields:

$$A_{X,t} \equiv \frac{b_X X_t}{\eta_X}, \qquad A_{P,t} \equiv \frac{b_P P_t}{\eta_P}. \tag{.7}$$

Using this conversion, the cost function—originally defined in terms of hours $(A_{X,t}, A_{P,t})$—can be expressed in terms of the proportional control parameters (b_X, b_P). The corresponding costs are given by:

$$C_X(b_X; X_t) = \theta_X \frac{b_X X_t}{\eta_X}, \qquad C_P(b_P; P_t) = \theta_P \frac{b_P P_t}{\eta_P}. \tag{.8}$$

Kelp dynamics. Using the relationship defined previously, the modified kelp dynamics under proportional fixed outplanting policy is given by:

$$X_{t+1} = \Big[k(X_t, P_t, R_t)\big(y(X_t, P_t, R_t) + \delta_X\big)\Big]\,(1 - \xi Z_t) + b_X \cdot X_t. \tag{.9}$$

Purple urchin dynamics. Similarly, under the proportional policy we subtract $b_P P_t$:

$$P_{t+1} = \max\Big\{P_{\text{exo}},\ \delta_P P_t + p_X(X_t, P_t) + p_D(X_t, P_t, R_t) - b_P \cdot P_t\Big\}. \tag{.10}$$

Value function under a fixed proportional policy. Let Π^F denote the Markov transition matrix on the discretized state grid induced by the proportional-transition rules in Equations (.9)–(.10) together with the distribution of the MHW shock. Let r^F be the corresponding one-period reward vector, computed at each grid

point using the proportional-cost expressions in (.8) (and the relevant stage objective). The infinite–horizon value vector V^F under this fixed policy satisfies the linear policy–evaluation system

$$\left(I - \beta \Pi^F\right) V^F \;=\; r^F, \tag{.11}$$

equivalently,

$$V^F \;=\; \left(I - \beta \Pi^F\right)^{-1} r^F. \tag{.12}$$

Here I is the $S \times S$ identity matrix, with S the total number of discrete states. After solving Equation (.11), we report the value evaluated at the degraded state of interest for Northern California kelp forests.

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 population 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 by[15]:

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

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.

[15]Normalization 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.

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 $\boldsymbol{\psi}$ be a vector of parameters of our interest (i.e., $\boldsymbol{\psi} = \{\gamma_P, \gamma_R, \mu_P, \mu_R, \ \ldots \ , \phi, \upsilon\}$) 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 $\boldsymbol{\psi}$ given the observed data is proportional to the product of prior and likelihood:

$$Pr(\boldsymbol{\psi}|Y_{obs}) \propto Pr(Y_{obs}|\boldsymbol{\psi}) \cdot Pr(\boldsymbol{\psi}), \tag{.13}$$

where $Pr(\boldsymbol{\psi})$ is the prior and $Pr(Y_{obs}|\boldsymbol{\psi})$ 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):

$$Pr(\boldsymbol{\psi}|Y_{obs}) = Pr(\boldsymbol{\psi}|s(Y_{obs})) \approx Pr(\boldsymbol{\psi}|d(s(Y_{obs}), s(\hat{Y})) \leq \Lambda). \tag{.14}$$

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:

$$D = \sum_{n=1}^{6} d_n \tag{.15}$$

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

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,[16] 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.

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.

[16] Using many summary statistics increases the likelihood of parameter rejections and computational complexity.

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.

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 (.15). 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.

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).

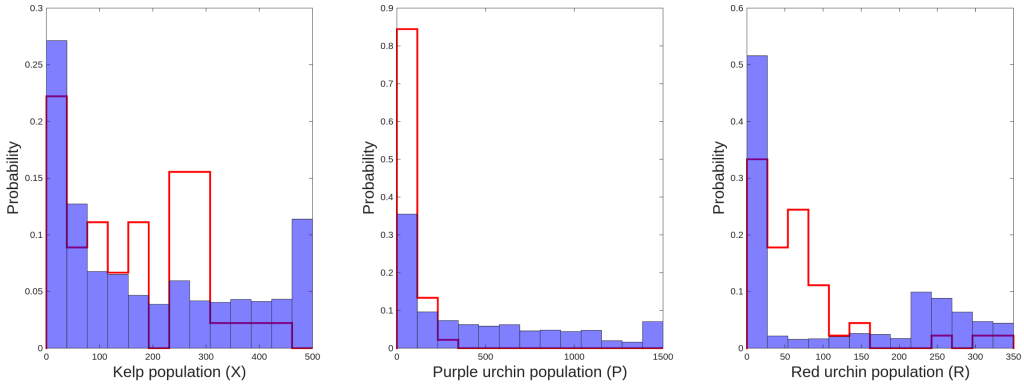


Figure C3: Comparison of probability density functions for kelp forest states based on sampled model parameters and observed data from Northern California. (Blue: computed long-run probability density based on sampled parameters; Red: observed long-run density). The computed density is obtained by applying the Markov transition matrix over a 50-year period and integrating over the other two state variables.

Extension: calculation of external recruitment

Furthermore, for the model modified with external recruitment, we provide the detailed procedure for identifying the external recruitment share parameter ζ and η_E. Following the estimated dispersal kernel for bull kelp spores by Burnett et al. (2024), which describes how far (i.e., the proportion of spores that travel away from their original recruitment site) bull kelp spores disperse, we assume that the remote source patch for external recruitment is located a few kilometers away (2–10 km). Based on this assumption, we examine candidate values of $\zeta \in \{0.05, 0.1, 0.15, 0.2, 0.25, 0.3\}$. For each candidate ζ, we search for the corresponding external recruitment parameter η_E that generates pre-MHW long-run probability density functions closely matching those observed for each species, together with the existing biological parameter values pre-estimated by ABC. As a result, we selected $\zeta = 0.05$ with $\eta_E = 23$, which provided the best fit, as shown in Figure C4.

Figure C4 presents results from the default parameterization with the newly introduced external recruitment setup and shows potentially improved fits to the long-run species density distributions, particularly for kelp and red urchin, relative to Figure C3, with the best-fitting values $\zeta = 0.05$ and $\eta_E = 23$. Under the symmetry assumption of the same ratio of loss of internal recruitment and influx of external recruitment, this improvement in fit is, biologically speaking, due to a reduced internal kelp recruitment share when the kelp population is high; because red urchin populations rely on kelp, this feedback decreases red urchin abundance and improves the fits. Furthermore, these values imply that, under density-dependent recruitment described

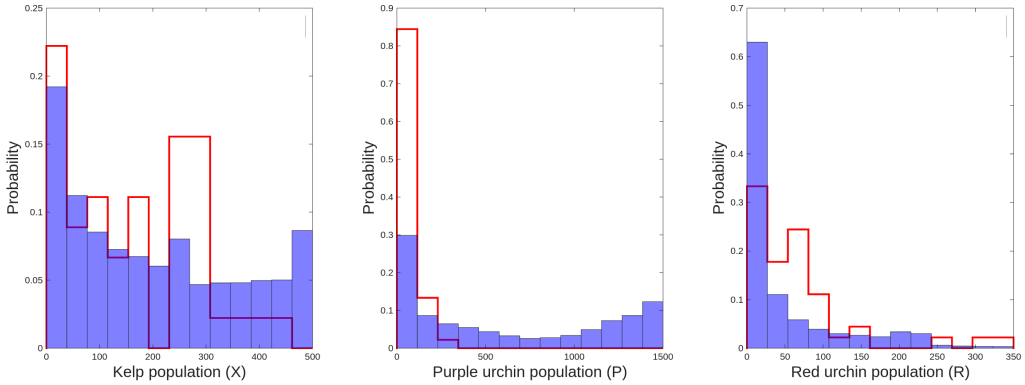


Figure C4: Estimated probability density functions for kelp forest states based on sampled model parameters, literature, and observed data from Northern California with external kelp spore recruitment ($\eta_E = 23$, and $\zeta = 0.05$) (Blue: computed long-run probability density based on sampled parameters; 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.

by the Beverton–Holt model, when the kelp population is low within the representative patch (possibly when urchin grazing is also low), external recruitment can play an important role in kelp recruitment even if spores arrive from sources several kilometers away. Conversely, when kelp density is sufficiently high, competition for settlement with internally produced spores strongly limits the contribution of external recruitment.

C.2 Marine heatwave (MHW) probability

The annual probability of MHW occurrence for each regime (α_i, for $i = \{low, medium, high\}$) 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.

Based on this dataset, the high MHW regime ($i = high$) 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.

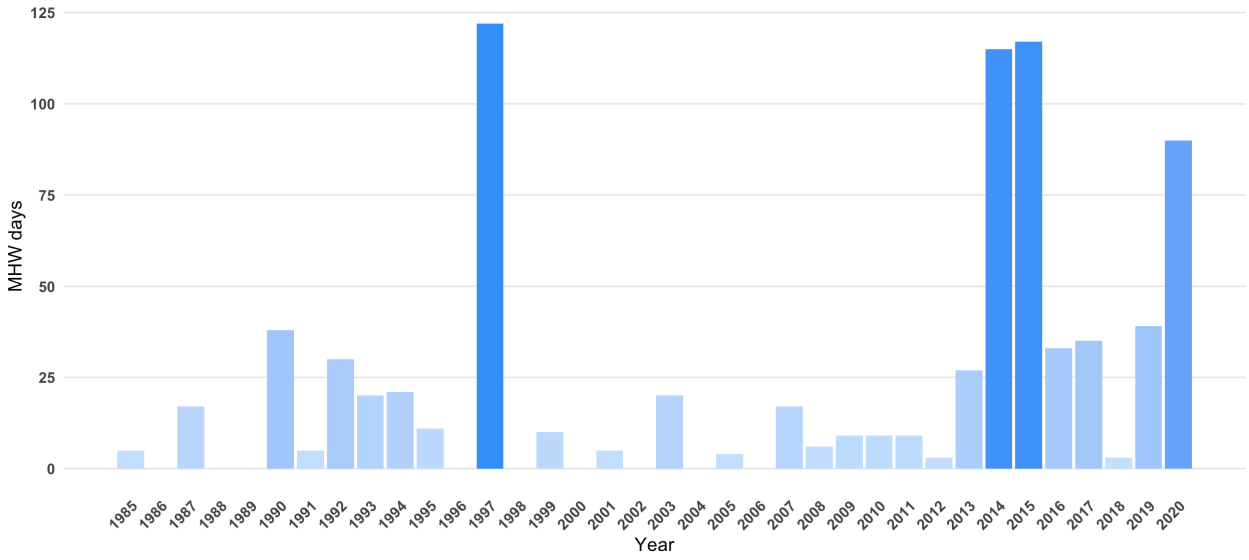


Figure C5: Historic marine heatwave (MHW) days from 1985 to 2020 for Mendocino and Sonoma Counties in Northern California. Data for 1985–2019 are from McPherson et al. (2021), with 2020 supplemented using records from NOAA Physical Sciences Laboratory (2025). For 2020, the MHW duration was approximately 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.

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 (Ward et al., 2022). 2) Total diving hours for restoration: 720 hrs (Ward et al., 2022). 3) Average weight of a purple urchin: 62 g/urchin (CDFW personal communication). Calculation: 14,292 kg / 62 g / 720 hrs = 320.16 urchins/hr.

3. $\max(A_X)$, maximum kelp outplanting effort range (120 hours): Based on Duggins et al. (2001), 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 $360m^2/1.5m^2/2.1$ 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), a 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 multipler 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 population 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/360 m^2): Based on the restoration report provided by Ward et al. (2022), purple urchin recolonization after removal is 17.6% of the pre-restoration population. With an average pre-MHW population of 300 urchins/360 m^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/360 m^2): Assuming the same recolonization rate and average population 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.

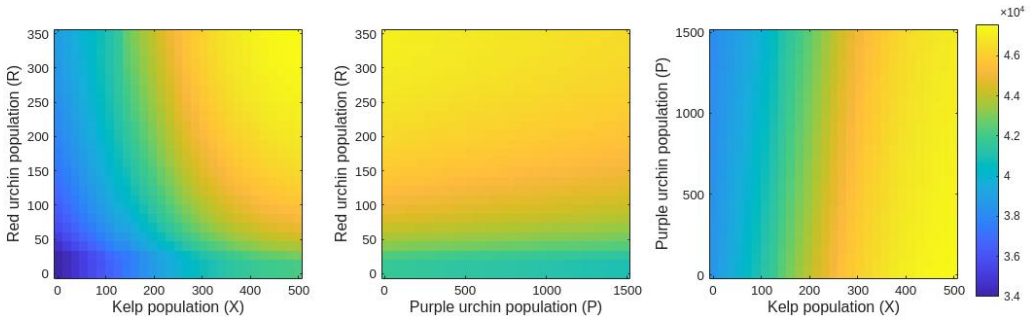
17. θ_X (Kelp outplanting cost per person hour, $42.5/hr): $382/15 hrs = $25.46/hr in 2002, adjusted for CPI, as per Carney et al. (2005).

18. θ_P (Purple urchin removal cost per diving hour, $102.6): Total cost for purple urchin removal: $65,490 / 720 hrs = $90.9/hr adjusted to $102.6/hr post-CPI adjustment, as per Ward et al. (2022).

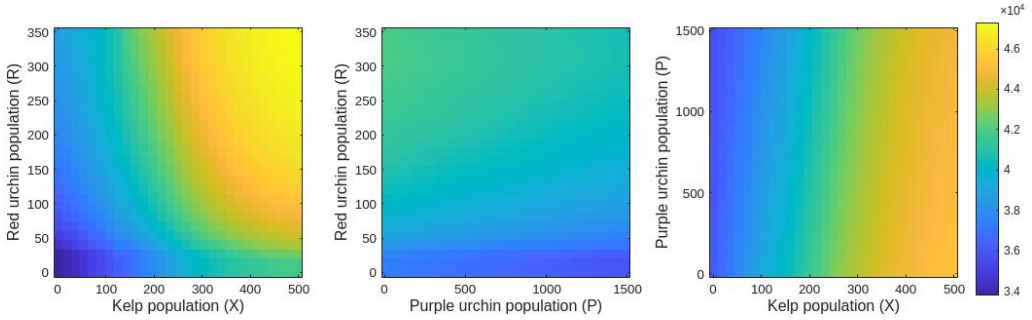
19. θ_H (Red urchin harvest cost per fishing trip, $16.21/patch): (1) Based on an interview with an urchin diver in Ft. Bragg. Minimum earning per trip is $500 with $150 for tender operation. Ward et al. (2022) states that divers participating in purple urchin removal paid $600 for a ship. We assume 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 Alternative scenario 1: Latent states

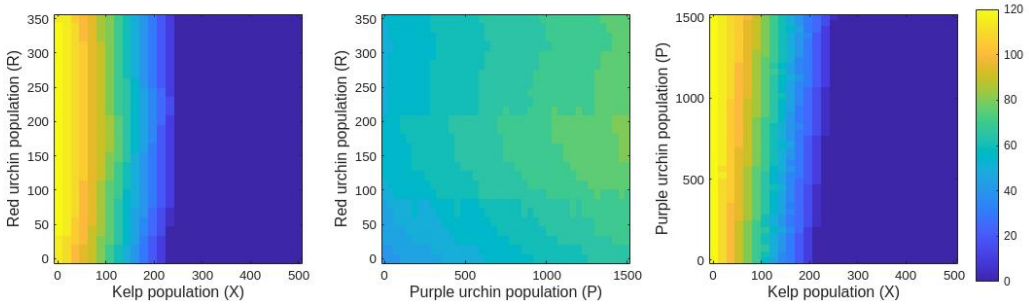


(a) Value function at the restored latent state

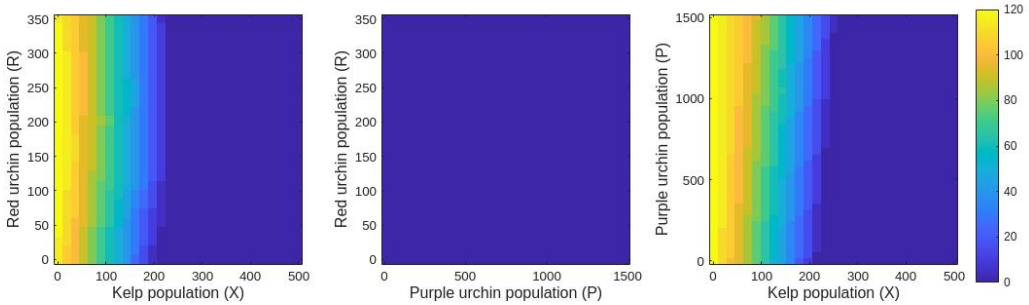


(b) Value function at the degraded latent state

Figure D6: Value functions given the baseline model except with restored (top) and degraded (bottom) latent states.

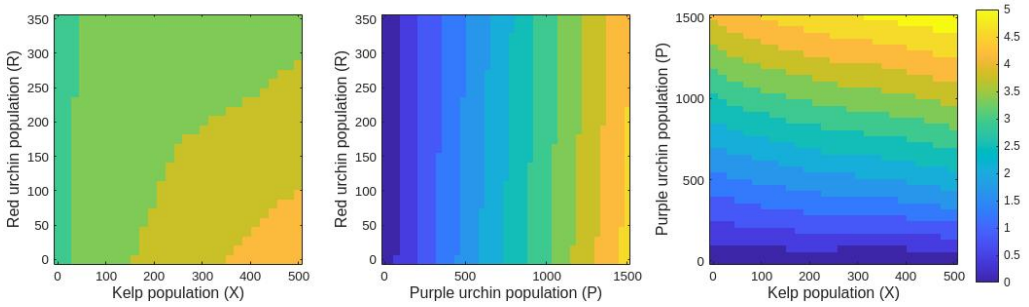


(a) Kelp outplanting intensity (degraded latent state)

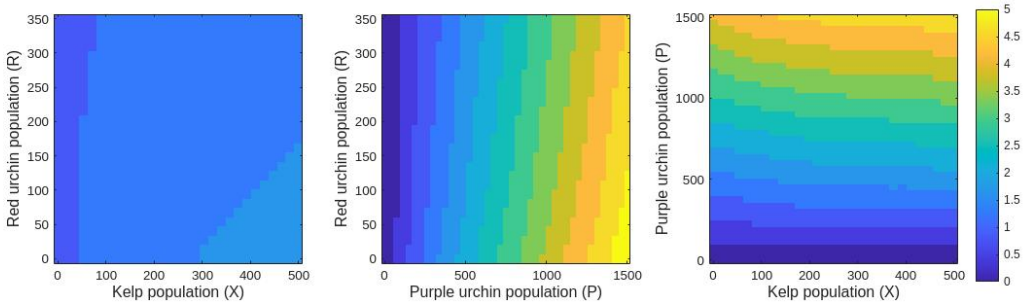


(b) Kelp outplanting intensity (restored latent state)

Figure D7: Optimal kelp outplanting intensity (labor hours) given the baseline model except with degraded (top) and restored (bottom) latent states.



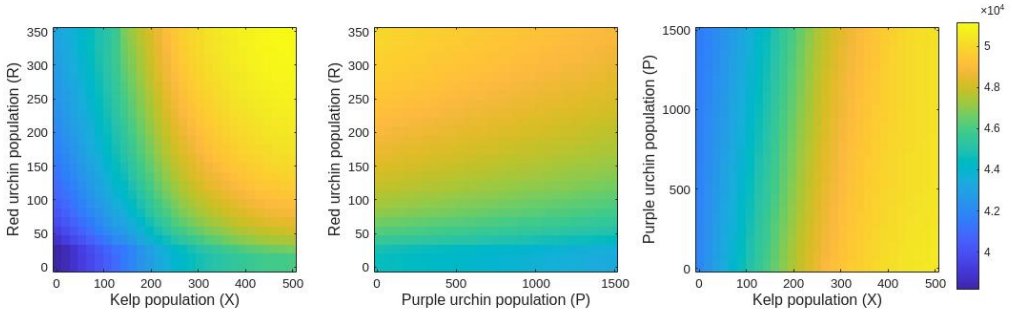
(a) Purple urchin removal intensity (degraded latent state)



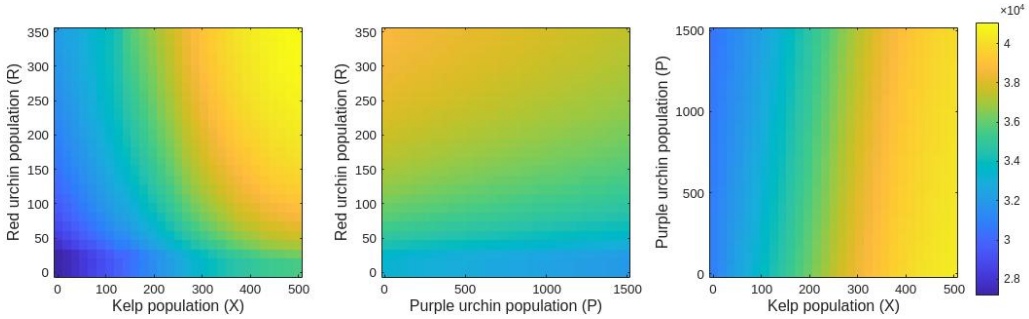
(b) Purple urchin removal intensity (restored latent state)

Figure D8: Optimal purple urchin removal intensity (labor hours) given the baseline model except with degraded (top) and restored (bottom) latent states.

D.2 Alternative scenario 2: MHW-frequency levels and uncertainty

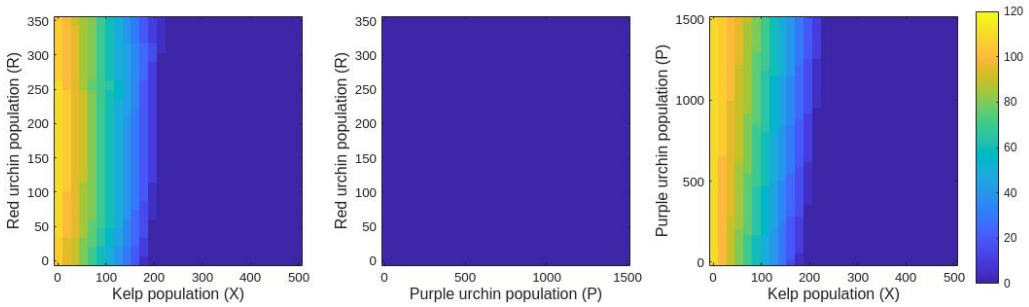


(a) Value function (low MHW regime)

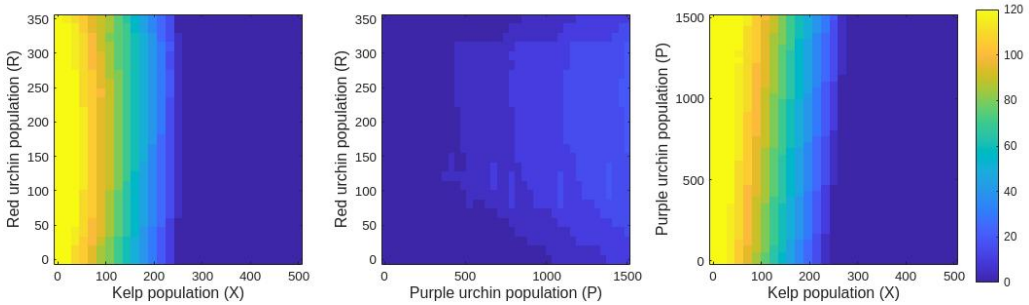


(b) Value function (high MHW regime)

Figure D9: Value functions given the baseline model except with low MHW regime (top) and high MHW regime (bottom).

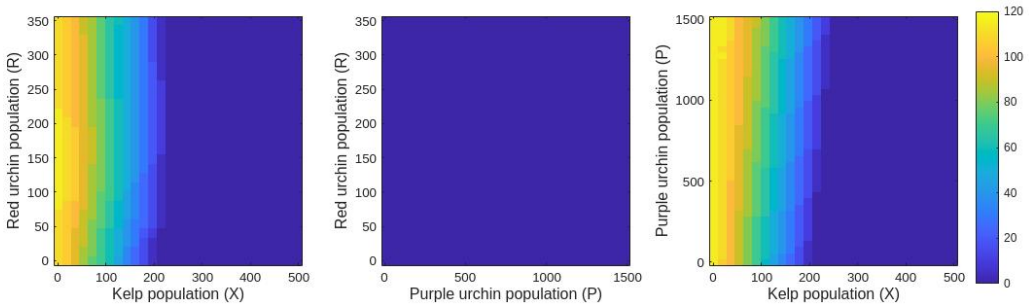


(a) Kelp outplanting intensity (low MHW regime)

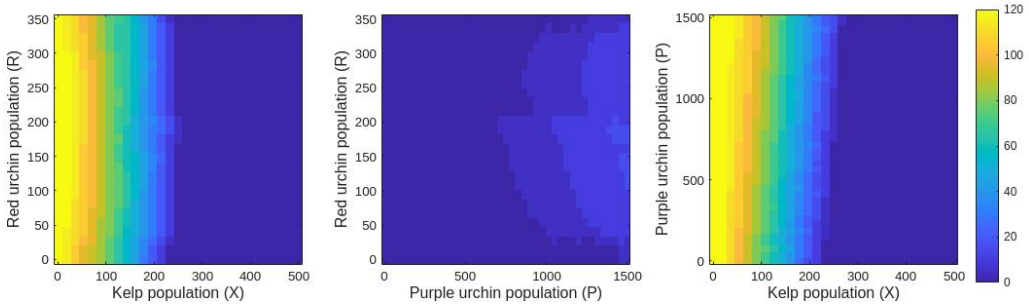


(b) Kelp outplanting intensity (high MHW regime)

Figure D10: Optimal kelp outplanting intensity (labor hours) given the baseline model except with low MHW regime (top) and high MHW regime (bottom).

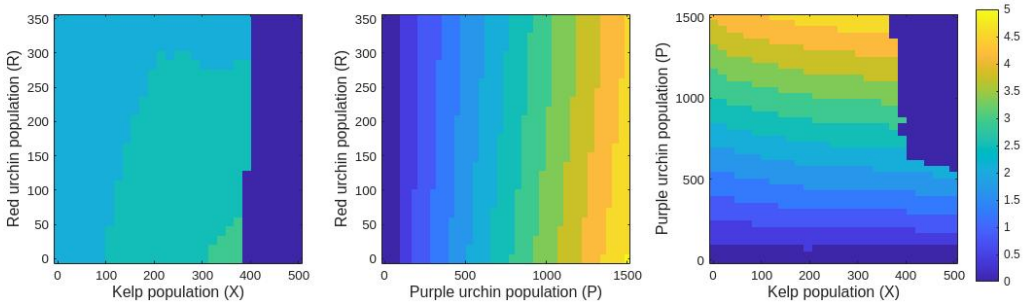


(a) Kelp outplanting intensity (unknown MHW; low-frequency-skewed belief)

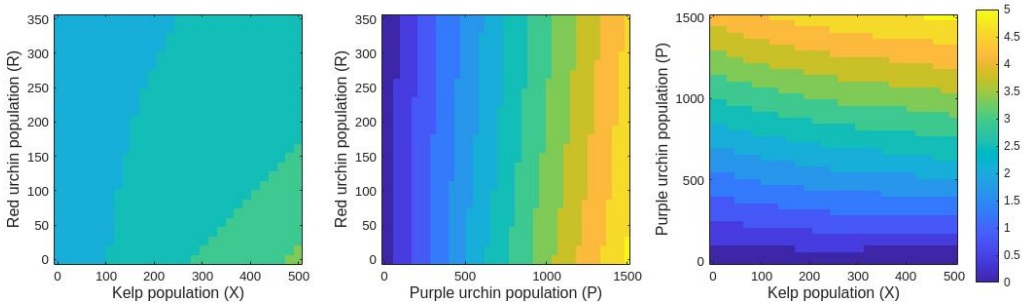


(b) Kelp outplanting intensity (unknown MHW; high-frequency-skewed belief)

Figure D11: Optimal kelp outplanting intensity (labor hours) given the baseline model except with unknown MHW regime with low-frequency-skewed belief (top) and high-frequency-skewed belief (bottom).

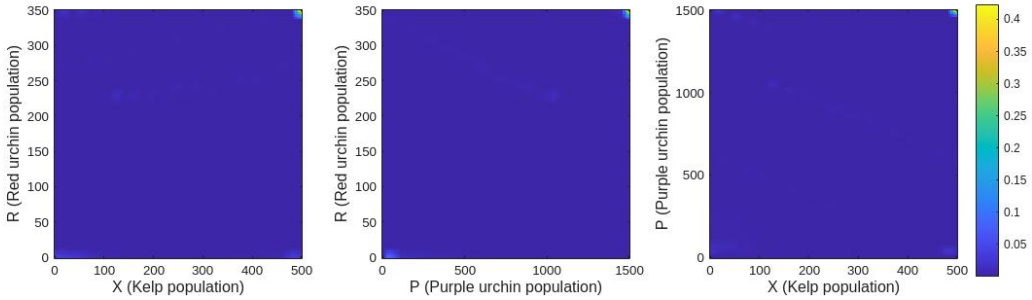


(a) Purple urchin removal intensity (low MHW regime)

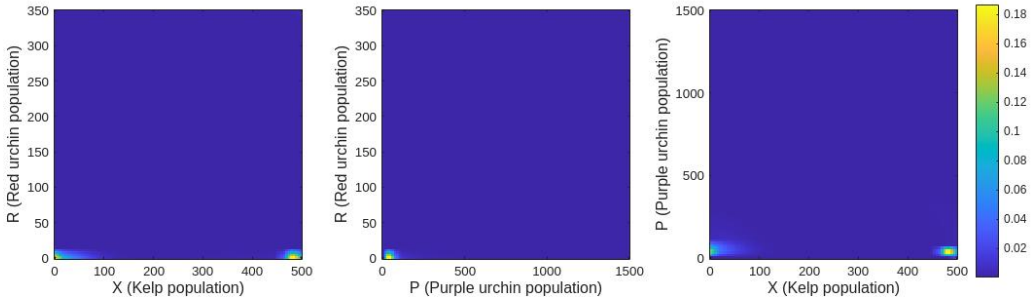


(**b**) Purple urchin removal intensity (high MHW regime)

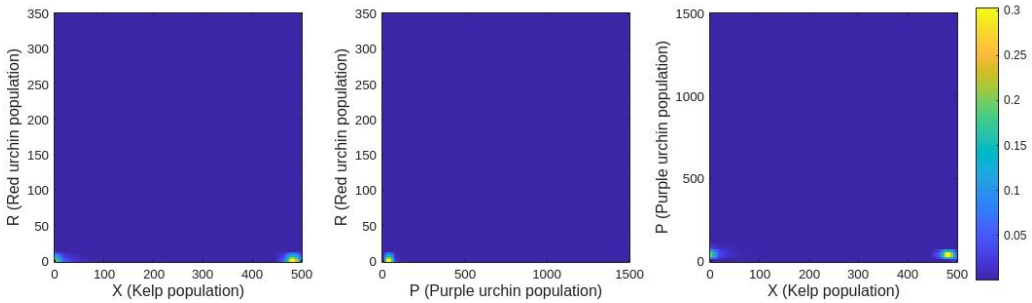
Figure D12: Optimal purple urchin removal intensity (labor hours) given the baseline model except with low MHW regime (top) and high MHW regime (bottom).



(a) No restoration (low MHW)

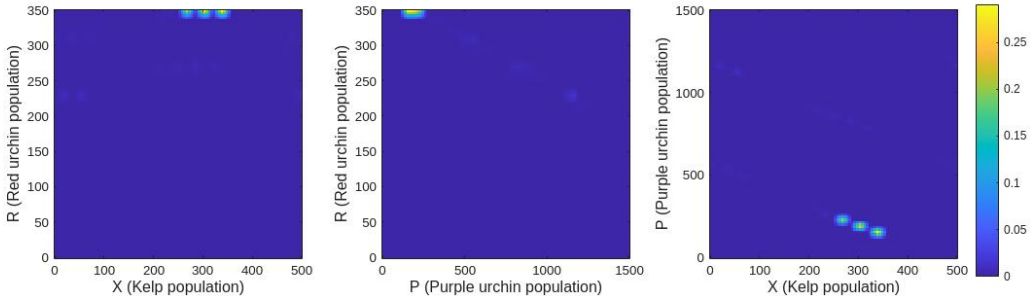


(b) No restoration (medium MHW)

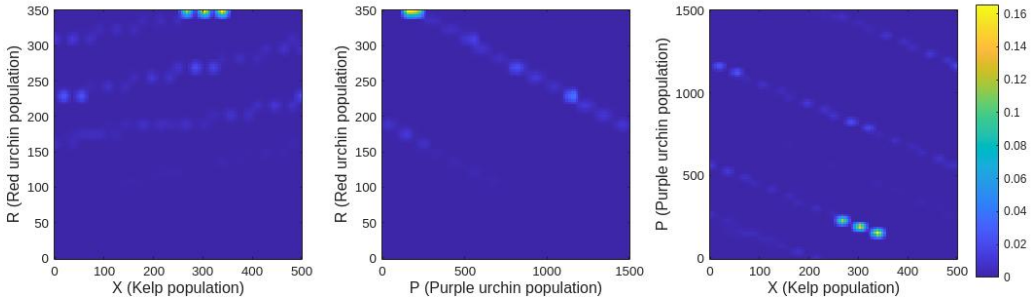


(c) No restoration (high MHW)

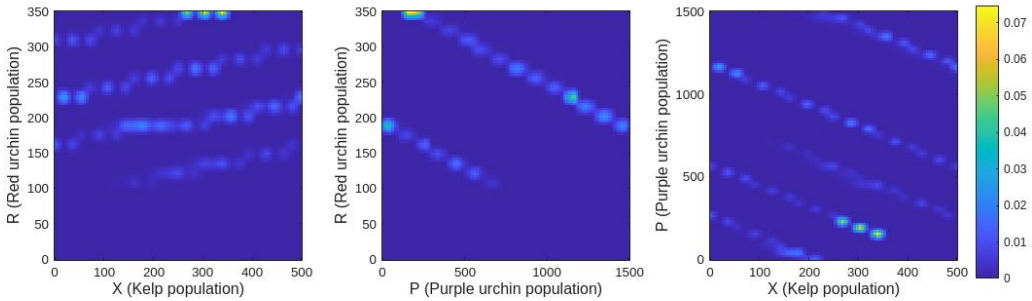
Figure D13: Bivariate long-run state distributions with no restoration given the baseline model except with low- (top), medium- (middle), and high-MHW (bottom) regimes.



(a) Optimal restoration (low MHW)

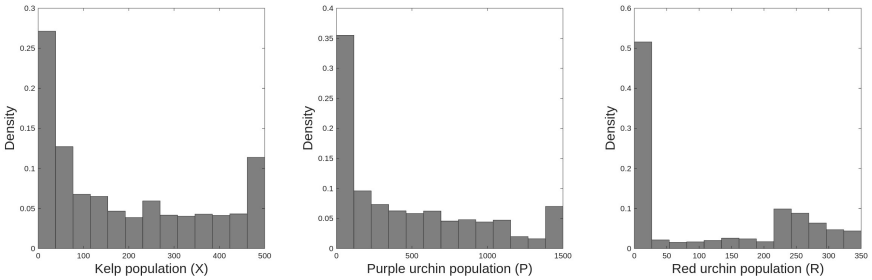


(b) Optimal restoration (medium MHW)

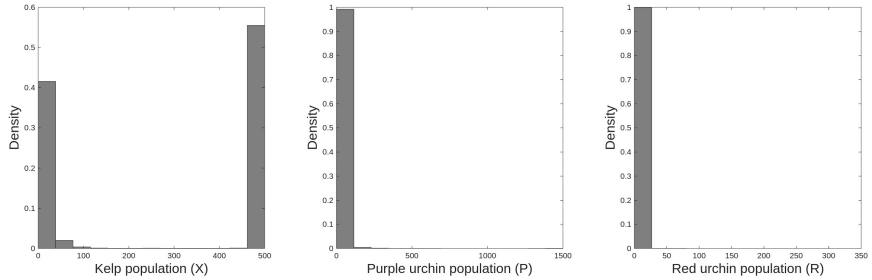


(c) Optimal restoration (high MHW)

Figure D14: Bivariate long-run state distributions with optimal restoration given the baseline model except with low- (top), medium- (middle), and high-MHW (bottom) regimes.

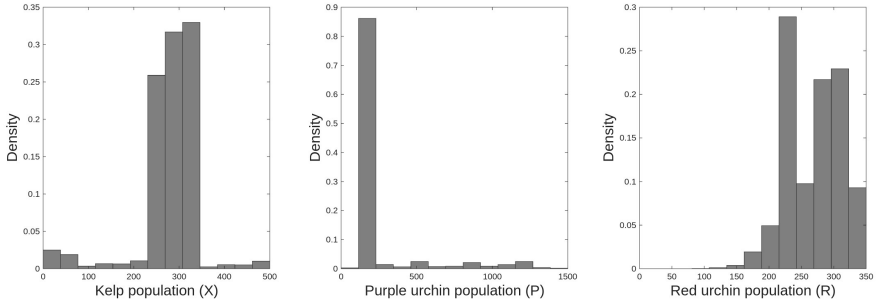


(a) No restoration (low MHW)

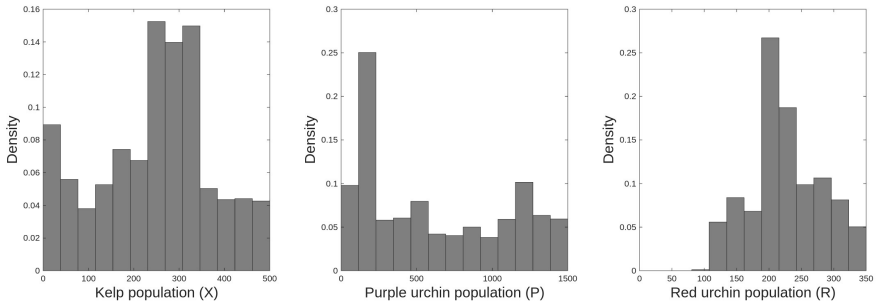


(b) No restoration (high MHW)

Figure D15: Probability density functions for state variable levels after 50 years under no restoration given the baseline model except with low MHW (top) and high MHW (bottom).



(a) Optimal restoration (low MHW)



(b) Optimal restoration (high MHW)

Figure D16: Probability density functions for state variable levels after 50 years under optimal restoration given the baseline model except with low-MHW (top) and high-MHW (bottom) regimes.

D.3 Alternative scenario 3: Kelp spore loss and exogenous recruitment

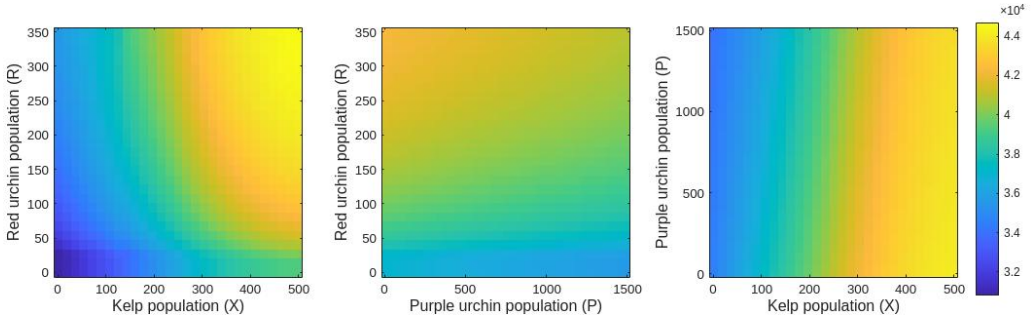
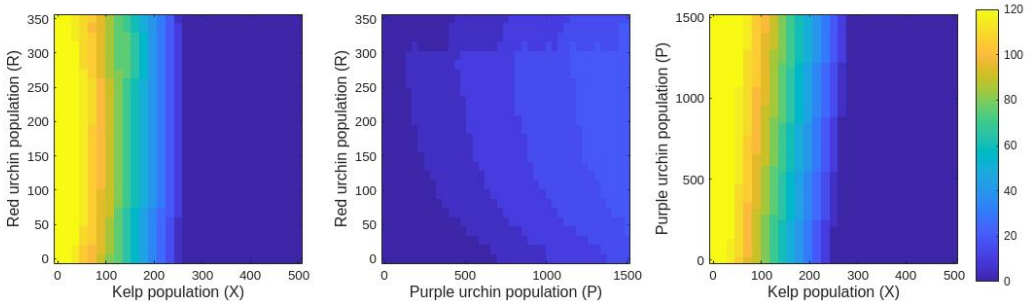
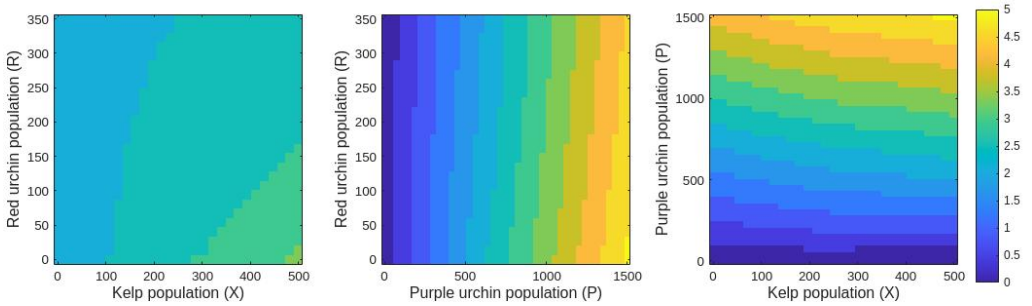


Figure D17: Value function given the baseline model except with external recruitment.



(a) Kelp outplanting intensity (external recruitment available)



(b) Purple urchin removal intensity (external recruitment available)

Figure D18: Optimal kelp outplanting (top) and purple urchin removal (bottom) intensity (labor hours) given the baseline model except with external recruitment.

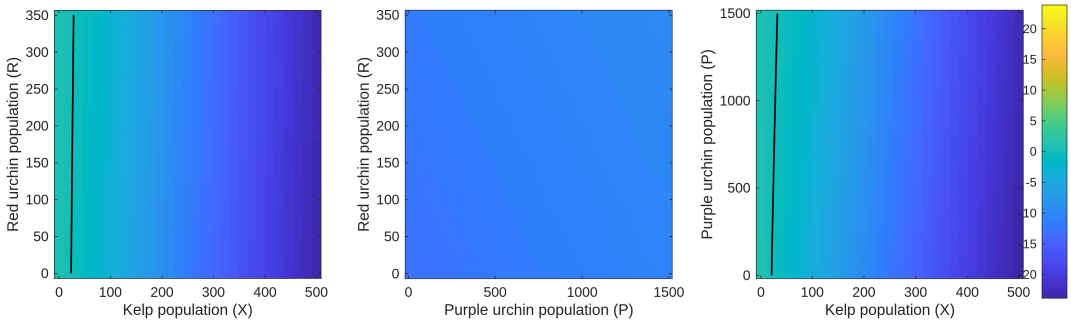


Figure D19: Difference in X_{t+1}, $\Delta X_{t+1} \equiv X^{\text{exo}}_{t+1} - X^{\text{no}}_{t+1}$, under external spore recruitment versus no external recruitment, evaluated under no restoration ($A_X = A_P = 0$). The black contour indicates $\Delta X_{t+1} = 0$.

D.4 Alternative scenario 4: Static proportional heuristic policy

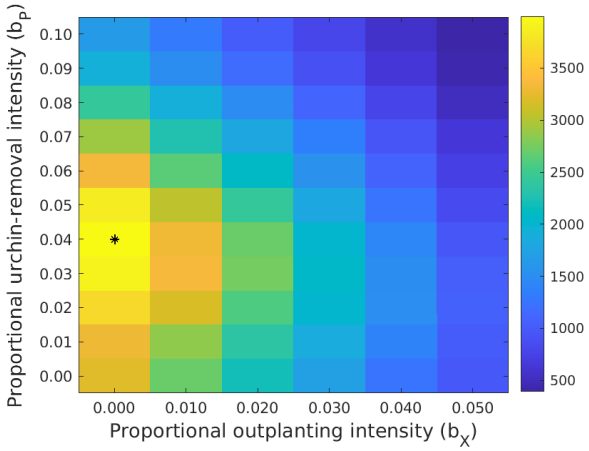


Figure D20: The value function showing the expected present value of long-run net benefits (color bar) under the baseline model but where restoration levels are determined by the static proportional policy. The starting state is given by the degraded scenario ($\mathcal{X} = 25$, $\mathcal{P} = 75$, $\mathcal{R} = 25$). The asterisk (*) indicates the levels of b_X and b_P that result in the maximum value across the static proportional policy alternatives. Outside the grid considered in Figure D20, values decline monotonically with increasing (b_X, b_P), eventually becoming negative (management costs outweigh any returns).

D.5 Alternative scenario 5: No return of the red urchin harvest

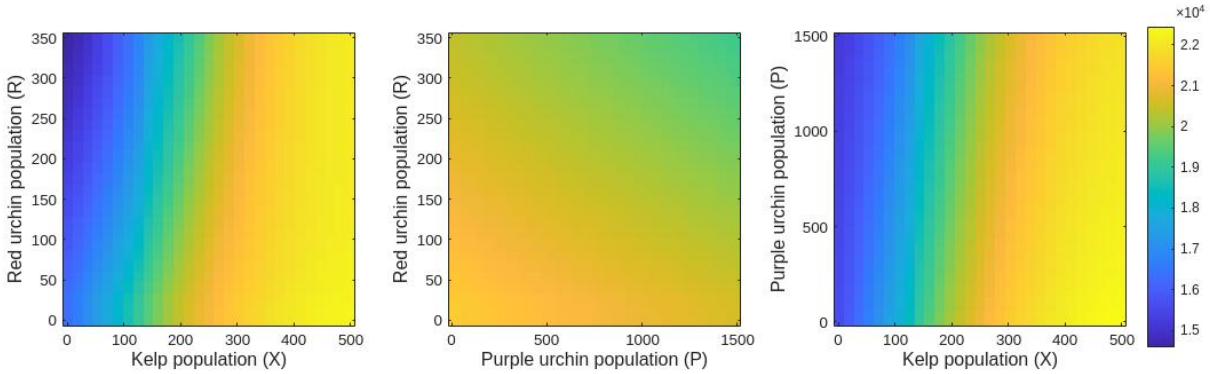
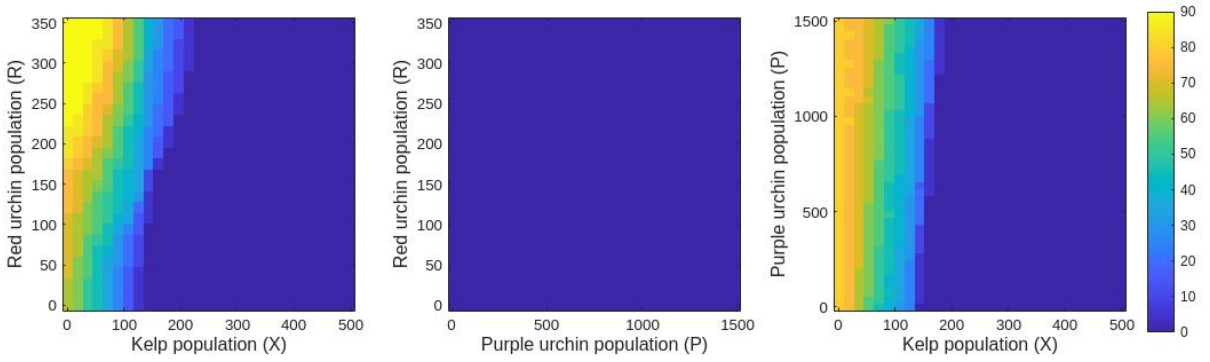
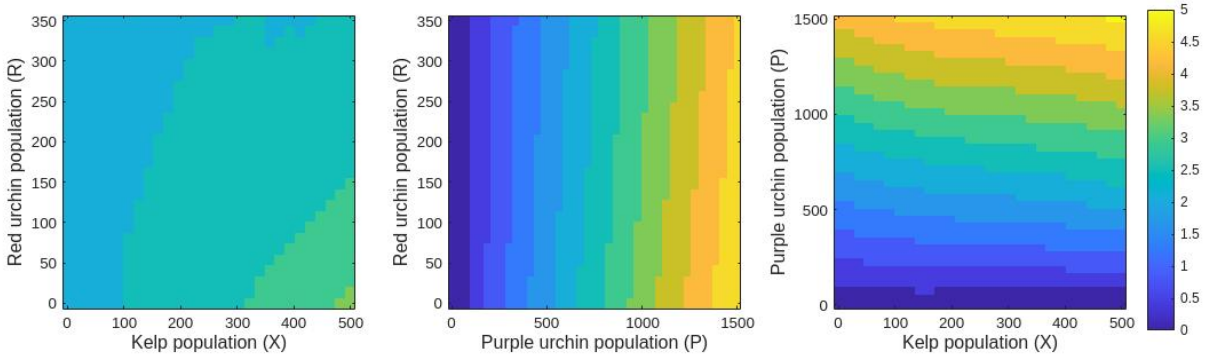


Figure D21: Value function given the baseline model except with no red urchin harvest.



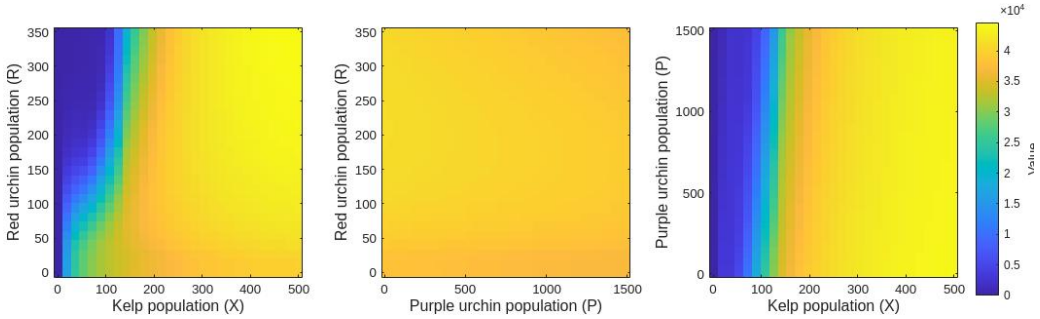
(a) Kelp outplanting intensity (no red urchin harvest benefit)



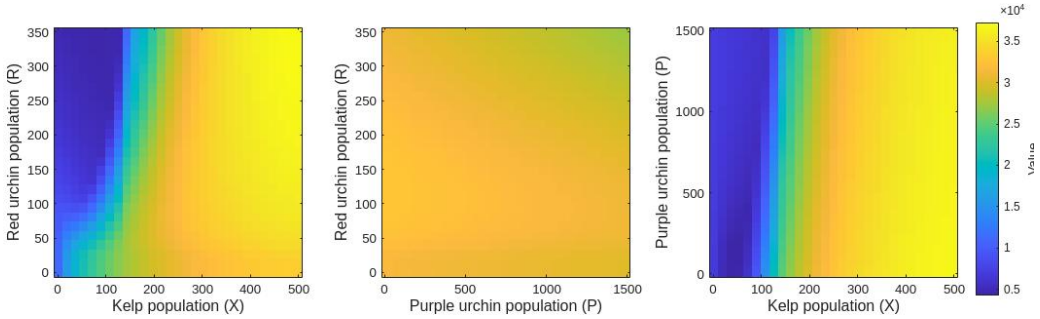
(b) Purple urchin removal intensity (no red urchin harvest benefit)

Figure D22: Optimal kelp outplanting (top) and purple urchin removal (bottom) intensity (labor hours) given the baseline model except with no red urchin harvest.

D.6 Alternative scenario 6: Restricted restoration

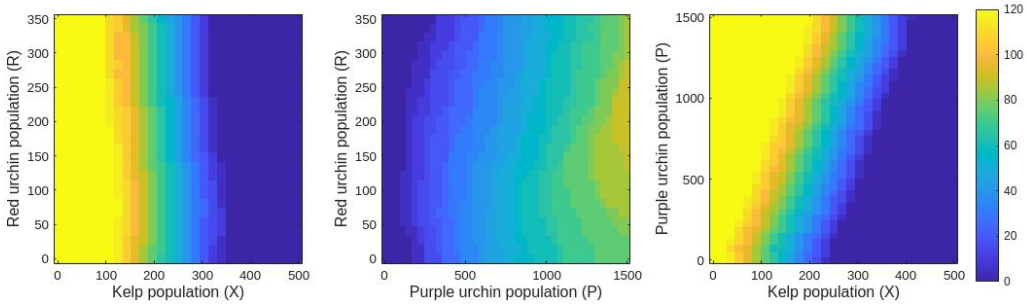


(a) Value function ($A_X = 0$)

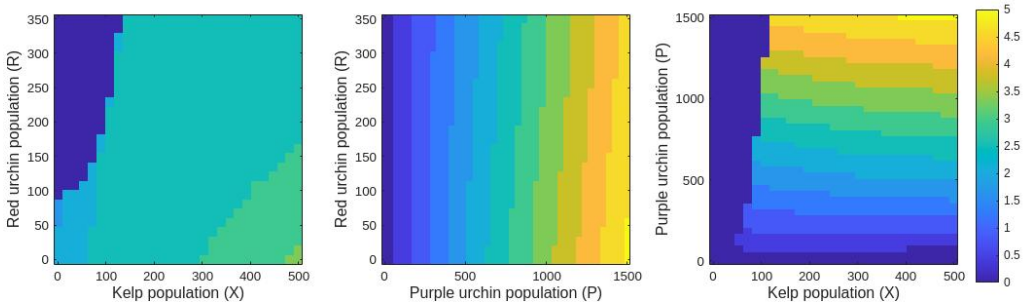


(b) Value function ($A_X = 0$; external recruitment)

Figure D23: Value functions given the baseline model except with restricted restoration ($A_X = 0$, top) and additionally with external recruitment (bottom).



(a) Kelp outplanting intensity ($A_P = 0$; external recruitment)



(b) Purple urchin removal intensity ($A_X = 0$; external recruitment)

Figure D24: Optimal restoration given the baseline model except with restricted tool choice and external kelp recruitment: outplanting only (top, $A_P = 0$) and removal only (bottom, $A_X = 0$).

Chapter 4

Dynamic Multiplier Effects of Commercial Fisheries in Local Fishing Economies

Abstract

Commercial fisheries are economically vital to many fishing communities, supporting employment across multiple non-fishing sectors of local economies, yet existing empirical evidence on how fishing sector performance affects local non-fishing employment and wages lacks long-run causal estimates. Research on the local economic effects of natural-resource sectors has concentrated overwhelmingly on non-renewable extraction such as oil and gas, typically summarizing effects as a single cumulative magnitude; for renewable resources, where sound management can sustain a perpetual flow of local economic activity, the dynamic path of multiplier effects has never been traced. This paper provides the first estimates of the explicit time path of local multiplier effects for a renewable resource using community-level data. Using data for approximately 150 Alaska fishing communities over 2000–2016, we employ a panel local projections framework with a shift-share instrument to trace causal cumulative impulse responses of local employment and wages to fishing revenue shocks horizon by horizon over five years, incorporating recent analytical bias correction methods to address the incidental parameter problem arising from a short time dimension with persistent outcomes. We find that fishing revenue shocks generate employment effects extending well beyond the contemporaneous period, following a non-monotonic trajectory that shifts toward positive cumulative spillovers at longer horizons. The dynamics differ across channels: resident permit-owner earnings generate a delayed but gradually building traded-sector response, while local landings produce more immediate but transient effects. Highly fishing-dependent communities experience amplified short-term volatility, including early-horizon negative effects consistent with labor crowding-out in small rural communities. These findings suggest that short-run policy evaluations may understate the cumulative returns to fisheries investments, and that the timing, channel, and community context of spillovers should inform fisheries management and community development strategies.

Keywords: fishing communities, dynamic multiplier effects, local projections, incidental parameter problem

4.1 Introduction

Commercial fisheries play a substantial role in coastal and regional economies across the United States, contributing to a broader marine economy that supports approximately 3.7 million jobs nationwide (NOAA,2024). Within this marine economy, commercial fishing and the seafood industry alone generated $5.9 billion in landings revenue and supported about 1.6 million jobs across the supply chain in 2022 (NOAA, 2022). Alaska stands out as the dominant fisheries producer, supplying roughly 60% of all U.S. domestic seafood landings. The Alaska seafood industry harvested 5.8 billion pounds of seafood in 2022, valued at $2.0 billion at the dock and $5.2 billion at the wholesale level after processing, and directly employed approximately 48,000 workers. Seafood processing is in fact the largest manufacturing sector in Alaska, accounting for 66% of the state's manufacturing employment (ASMI, 2024).

Under National Standard 8 of the Magnuson-Stevens Fishery Conservation and Management Act, agencies are required to consider the economic impacts of management decisions on fishing communities and to minimize adverse effects on their sustained participation (Clay and Olson, 2008). This mandate has grown more consequential as coastal communities face increasing competition for limited offshore space from aquaculture, tourism, offshore wind energy, and other marine industries (FAO, 2024), requiring managers to evaluate the full direct, indirect, and induced economic impacts of fishing relative to competing uses (Jacobsen et al., 2014). Without credible long-run evidence, however, decisions may rely on shorter time horizons and incomplete indicators, increasing the risk of myopic resource management that affects the economic welfare of fishing communities. Understanding how impacts from the fishing sector spread to the local economy over multiple years, and how long these effects persist, would allow agencies to better anticipate the consequences of policy changes and to design more durable policies around stock rebuilding, conservation, and restoration.

Despite this economic significance and the policy demands surrounding it, a fundamental question remains underexplored: when a fishing community experiences a change in harvest, i.e., fishing revenues, how does that change propagate to non-fishing sectors of the local economy over time? Although policymakers routinely assume that fishing performance is interwoven with employment and wages in the broader local economy, this connection has rarely been tested empirically (Watson et al., 2021), and its long-run dynamics have never been estimated. This question has grown increasingly urgent as climate-induced changes in the ocean compound existing pressures on fishing communities. Shifts in the distribution and productivity of

fish stocks, together with overexploitation, ocean acidification, and the globalization of the seafood trade, can destabilize income and employment in fishing-dependent communities (Free et al., 2019; Rogers et al., 2019; Pinsky et al., 2018; Anderson et al., 2018). If the relationship between fisheries and local economic outcomes is persistent, then such disruptions may have lasting consequences well beyond the harvesting sector.

The existing empirical literature on fisheries and local economies is limited in both scope and methodology. Much of what is known comes from input-output and simulation-based models. In the Alaska context, Seung and Waters (2009) quantify linkage effects using a social accounting matrix model, and Seung and Waters (2010) employ a computable general equilibrium framework to estimate multiplier effects of the seafood industry. While useful for characterizing economic linkages, such frameworks rely on fixed technical coefficients or assumed elasticities, which limits their ability to capture time-varying dynamics in sectoral linkage structures. Among the few econometric studies, Seung (2008) estimates dynamic employment impacts from seafood processing using a vector error correction model for two fishery-dependent regions in Alaska, finding relatively small effects on non-seafood employment. However, the impulse response functions (IRFs) lack a causal interpretation, as identification relies on variable ordering assumptions rather than economically justified exogeneity; the study examines only processing employment shocks rather than total fishing revenue, missing the permit-owner earnings channel later identified by Watson et al. (2021); and the analysis is limited to two aggregated regions, precluding the use of cross-sectional variation.

On the causal side, Watson et al. (2021) provide the first empirical evidence of fisheries multiplier effects, finding that a dollar increase in fishery earnings raises total resident income by approximately 1.54 dollars, primarily through resident permit-owner earnings rather than the location of fish landings. Their estimates, however, are limited to contemporaneous, within-year effects. A contemporaneous multiplier is silent on whether induced spending effects accumulate over multiple years or dissipate quickly as seasonal workers depart and non-resident owners redirect earnings elsewhere. If the cumulative multiplier over a longer horizon differs substantially from the contemporaneous estimate, policy evaluations anchored to short-run evidence may systematically misjudge the long-run returns to fishery conservation, stock rebuilding, or quota reallocation.While Watson et al. (2021) establish the causal channel and magnitude of fisheries multiplier effects, the temporal structure of these effects remains an open question that this paper seeks to answer.

A substantial macroeconomic literature on fiscal policy shows that the dynamic effects of consolidations and spending shocks differ markedly from their contemporaneous impacts, with multipliers varying by time

horizon and by the state of the economy (Auerbach and Gorodnichenko, 2012; Ramey and Zubairy, 2018; Jordà and Taylor, 2016). The same horizon dependence is increasingly recognised in the local labor markets literature. Moretti (2010) estimates that each additional traded-sector job generates roughly 1.6 non-traded jobs in the long run, with the magnitude governed by local labor supply elasticity, while Van Dijk (2018) notes that estimated multipliers can differ with the length of the observation window, potentially rising as migration has more time to operate or falling as crowding out grows. Closest to the present setting, Allcott and Keniston (2018) examine the case of natural resource sectors, identifying oil and gas booms from the interaction of resource endowment with national price shocks, a strategy closely related to the shift-share instrument used here, and find that these booms raise local wages without crowding out manufacturing, with no evidence of a resource curse in the modern United States.

In this paper, we ask two specific questions. First, do fishing revenue shocks generate cumulative employment and wage effects that persist beyond the contemporaneous period, and if so, how quickly do they build, when do they peak, and how long do they last? Second, how do these dynamic multiplier effects vary across sectoral channels, sources of fishing revenue, and levels of community fishing dependency?

We address these questions by providing the first estimates of long-run dynamic multiplier effects from commercial fisheries on local economies. Research on the economic effects of natural-resource sectors has concentrated overwhelmingly on oil and gas (Allcott and Keniston, 2018; Arezki et al., 2017), while the evidence for renewable resources remains thin, with Watson et al. (2021) the leading exception. This gap matters: unlike depletable oil and gas, a renewable resource can sustain a perpetual flow of local economic activity under sound management, so the policy-relevant question becomes the long-run return to conserving or rebuilding the stock. Beyond their focus on non-renewable resources, the local estimates that do exist typically identify effects by comparing outcomes across widely spaced Census waves (i.e., long-difference estimation), summarizing the response as a single cumulative magnitude rather than a path; where richer dynamics are recovered, they come from country-level resource-windfall studies or event-study treatments of non-renewable boom-bust episodes, not from a local multiplier traced horizon by horizon from a continuous revenue shock. The one available local renewable-resource estimate, in turn, is contemporaneous. Building on the identification strategy of Watson et al. (2021), we employ a shift-share instrument leveraging exogenous variation in seafood prices and fish stocks to trace cumulative impulse responses of local employment and wages horizon by horizon over multiple years.

Our contributions are twofold. First, substantively, we situate commercial fisheries, a renewable resource,

within a literature dominated by non-renewable extraction, and instead of collapsing the response into a single number we trace the local multiplier horizon by horizon within a dynamic local-projection framework, estimating how quickly multiplier effects build, when they peak, and how long they persist beyond the contemporaneous period. To our knowledge, this is the first paper to estimate an explicit time path of local multiplier effects for a renewable resource using micro-level community data, giving policymakers the dynamic evidence needed to evaluate the long-run returns to fishery conservation and restoration.

Second, methodologically, we are among the first to bring a rapidly evolving strand of dynamic panel estimation to bear on a resource-economics question. Community-level data of this kind are rich in the cross-section but short in the time dimension, with many communities observed over relatively few years. In such a large-N, short-T panel with highly persistent outcomes, the fixed-effects estimator underlying standard local projections carries an incidental parameters bias (i.e., *Nickell* bias of dynamic panels) that does not vanish as the number of communities grows, distorting the estimated impulse responses and the long-run multipliers cumulated from them. We address this head-on, incorporating a recently developed analytical bias correction and verifying, through Monte Carlo simulations calibrated to replicate the dimensions and persistence of our own data, that the corrected estimator recovers the true response paths.

4.2 Methods and Data

This section presents our empirical framework. Section 4.2.1 describes the local projections model with instrumental variables used to estimate causal impulse response functions. Section 4.2.2 discusses the bias correction procedure for short panels with persistent outcomes. Section 4.2.3 validates our approach through Monte Carlo simulations. Section 4.2.4 describes the data.

4.2.1 Local Projections with Instrumental Variables (LP-IV)

We estimate long-run multiplier effects from the fishing sector on local economies using the local projections (LPs) method proposed by Jordà (2005). LPs offer several advantages over Vector Autoregressive (VAR) models for our setting. First, LPs rely on single-equation estimation, making them more robust to misspecification than VARs, which require a fully specified system. Second, LPs estimate impulse responses at each horizon separately, accommodating nonlinearity more naturally. Third, LPs provide a straightforward framework for estimating cumulative responses and multiplier effects in panel structures (Jordà, 2005; Jordà and

Taylor, 2025). Finally, LPs readily accommodate instrumental variable estimation for causal identification (Jordà et al., 2015; Ramey and Zubairy, 2018; Jordà et al., 2020).

For a panel where community i is observed at year t, the basic LPs model is:

$$y_{i,t+h} - y_{i,t-1} \coloneqq \Delta^h y_{i,t+h} = \mu_i^h + \lambda_t^h + \beta^h s_{i,t} + \upsilon_{i,t+h}, \quad h = 0, 1, 2, \ldots, H, \tag{4.1}$$

where μ_i^h is a community fixed effect, λ_t^h is a time fixed effect, and $y_{i,t+h} - y_{i,t-1} \coloneqq \Delta^h y_{i,t}$ is the long difference of the outcome over each h-period horizon. The outcome $y_{i,t}$ represents logged local economic variables such as total wage income and employment in non-fishing sectors. The shock variable $s_{i,t}$ is the log difference of fishing revenue, $s_{i,t} = \log(rev_{i,t}) - \log(rev_{i,t-1})$, where $rev_{i,t}$ denotes fishing revenue in community i in year t. Following Watson et al. (2021) and standard practice in the macroeconomic literature, this log-difference specification allows the estimated coefficient β^h to be interpreted as an elasticity, i.e., the percent change in the outcome from a 1% change in fishing revenue.

The cumulative impulse response can be captured using either a level specification or a long-difference specification in the outcome. However, we adopt the long-difference specification for three reasons: it directly captures cumulative multiplier effects, it provides robustness to unit root concerns (Jordà, 2023; Piger and Stockwell, 2023, 2025), and it mitigates finite-sample bias arising from persistent outcomes, as we discuss in detail in Section 4.2.2.

The residuals of a local projection typically exhibit a moving-average structure. While this does not affect the consistency of the estimated response, the induced serial correlation distorts standard errors and inference (Jordà, 2005, 2023). In a time-series setting, two common remedies are the Newey-West HAC estimator and lag augmentation paired with heteroskedasticity-robust (Huber-White) standard errors, as proposed by Olea and Plagborg-Møller (2021); the latter delivers valid inference under both stationary and highly persistent data and at long horizons.

However, our setting is a panel with large N and short T, where the appropriate inference tool differs. We report standard errors clustered by community, suitable when N is large relative to T (Jordà and Taylor, 2025; Jordà, 2023).[1] In this regime the asymptotic distribution of the response coefficient is governed by the cross-sectional dimension, so the near-unit-root distortions that motivate lag augmentation for inference do

[1] In a panel, the natural HAC analog is the Driscoll-Kraay estimator, whose asymptotic justification requires $T \to \infty$ (with N fixed or growing more slowly than T) and is therefore ill-suited to our short-T panel; clustering by community is instead justified under $N \to \infty$ with T fixed (Jordà and Taylor, 2025; Jordà, 2023).

not arise (Jordà and Taylor, 2025).

We nonetheless include lags of the outcome and the shock, and in particular the lagged outcome, for two reasons. Because First, they control for predetermined dynamics and the lingering effects of previous-period shocks, as is standard in applied local projection work (e.g., Romer and Romer, 2017; Mian et al., 2017; Baron et al., 2021). Second, and critically, the lagged outcome is important for the small-sample bias correction of Herbst and Johannsen (2024).[2] The baseline local projection in Equation (4.1) is thus extended as follows:[3]

$$\Delta^h y_{i,t+h} = \mu_i^h + \lambda_t^h + \beta^h s_{i,t} + \sum_{l=1}^{p} \gamma_l^h \Delta y_{i,t-l} + \sum_{l=1}^{q} \delta_l^h s_{i,t-l} + \upsilon_{i,t+h}, \quad h = 0, 1, 2, \ldots, H, \tag{4.2}$$

where p and q denote the lag lengths for the outcome and shock variables, selected via Monte Carlo simulations as detailed in Section 4.2.3.

Community-level fishing performance, $s_{i,t}$, is potentially endogenous because fishing activity and the local economy may be simultaneously determined. For instance, a community's harvesting and delivery decisions respond to the same local economic conditions that also drive our outcomes, creating simultaneity, as Watson et al. (2021) note. To address potential endogeneity of community-level fishing performance, we instrument $s_{i,t}$ using the shift-share instrumental variable strategy developed and validated by Watson et al. (2021):

$$z_{i,t} = \sum_j w_{i,j,0} \Delta x_{j,t}, \tag{4.3}$$

where $\Delta x_{j,t}$ is the fishery-wide growth rate of earnings from fishery j across all communities (the "shift"), and $w_{ij,0}$ is the pre-sample share of community i's commercial fishing earnings attributable to fishery j (the "share"). The instrument exploits exogenous variation in state- and national-level seafood prices and biological fish stocks, weighted by each community's historical fishery composition, using three pre-sample periods of 1998–2000. Individual communities have negligible influence over these price and stock movements, so the instrument isolates plausibly exogenous variation in fishing performance at the community

[2] We show in Appendix D that including lagged controls, in particular the lagged outcome, is critical for this bias correction.

[3] Most of the local projections literature adopts a parsimonious specification with few controls beyond lags; Watson et al. (2021), for instance, estimates without additional controls. Other community-level variables we considered, such as population, may introduce a "bad control" problem: poor fishing performance can trigger out-migration, which in turn may affects local employment and economic growth (Himes-Cornell and Hoelting, 2015; Carothers, 2010). We therefore keep the specification parsimonious, augmenting the baseline local projection only as needed for identification.

level.[4]

In the LP-IV framework, valid instrumentation requires an additional condition beyond the standard relevance and exclusion restrictions: lead-lag exogeneity (Stock and Watson, 2018). This condition requires that the instrument $z_{i,t}$ be uncorrelated with structural shocks to the local economy at all other time periods, not only contemporaneously. In our setting, this is plausible because the shift component is determined by forces external to any individual community. Furthermore, fixing the share weights $w_{i,j,0}$ at their pre-sample averages strengthens this condition: if shares were updated contemporaneously, local economic shocks could alter a community's fishery portfolio composition, feeding back into future values of the instrument and potentially violating lead-lag exogeneity. By holding shares fixed at pre-sample values, the instrument varies over time only through the exogenous shift component, eliminating this feedback channel (Watson et al., 2021; Borusyak et al., 2024).

4.2.2 Approximate Bias Correction for Dynamic Panel Local Projections

Our panel consists of a relatively large number of cross-sectional units ($N \approx 150$ communities) and a short time dimension ($T = 17$). In such settings, the incidental parameter problem (IPP) arises from the inclusion of community-specific fixed effects, whose estimation error does not vanish under fixed T even as N grows large. In dynamic panel models that include a lagged dependent variable, this problem manifests as the well-known Nickell bias (Nickell, 1981): the within transformation that removes the fixed effects induces a correlation between the demeaned lagged outcome and the demeaned error, biasing the estimates under fixed T regardless of how large N becomes (Alvarez and Arellano, 2003; Klosin, 2024). As we discuss below, our specification includes lagged outcome controls and is therefore subject to this bias.

Klosin (2024) challenges two common justifications for excluding the lagged outcome from fixed effects panel regressions: the concern that including $y_{i,t-1}$ introduces Nickell bias, and the belief that lagged outcomes need not be controlled for when treatment is randomly assigned conditional on fixed effects. The study shows that omitting the lagged outcome instead induces a separate bias, which the research terms *dynamic bias*, that typically exceeds the Nickell bias it is meant to avoid. The two biases differ in how they reach the treatment effect estimate: Nickell bias falls primarily on the autoregressive coefficient and contaminates the treatment effect only indirectly, whereas dynamic bias arises when the omitted lagged

[4] Watson et al. (2021) provide a detailed validation of the shift-share instrument, demonstrating that it satisfies both the relevance and exclusion conditions.

outcome is absorbed into the error, and the within transformation mixes errors across periods so that the demeaned error becomes correlated with the treatment itself, biasing it directly. This direct channel explains why dynamic bias dominates across the parameter ranges in Klosin (2024), and why it persists even when treatment is exogenous conditional on fixed effects.

In a related context tailored to LPs, Herbst and Johannsen (2021, 2024) analytically show that including outcome lags truncates the regression error from an MA(∞) to a finite-order MA process, so that the approximate bias depends only on impulse response coefficients at horizons up to h rather than on the entire IRF path. This truncation does not by itself guarantee a smaller bias, since the relative magnitude of the bias with and without lag controls depends on the underlying data-generating process (Herbst and Johannsen, 2024). Its value lies elsewhere: the bias becomes a finite sum of already-estimated quantities, which is what makes the analytical correction described below feasible. Two additional results indicate that lagged outcome controls also reduce the bias in magnitude in our dynamic panel setting. Klosin (2024) shows that the dynamic bias from omitting the lagged outcome typically exceeds the Nickell bias its inclusion introduces, and our Monte Carlo simulations calibrated to our data setup (Appendix E) confirm that specifications with lagged outcome controls exhibit substantially smaller bias than those without. Importantly, this literature demonstrates that even when the shock of interest is exogenous or identified through instrumental variables, the estimates may still suffer from bias under high persistence of the outcome variable and a finite time dimension. We provide a formal justification in Appendix A.

With the inclusion of outcome lags to avoid the larger dynamic bias that arises from their exclusion (Klosin, 2024), the remaining bias, as characterized by Nickell (1981), is primarily driven by two factors: the persistence of the outcome variable, captured by its autoregressive parameter ρ, and the length of the time dimension T. The bias is of serious concern only when both conditions hold simultaneously: high persistence and short T. In our setting, the long-difference specification further mitigates the persistence driver. Specifically, Piger and Stockwell (2025) show that differencing cancels out much of the covariance between the outcome and error terms, with bias reduction particularly pronounced even when the autoregressive process is near unit root. Furthermore, the yearly data frequency may additionally reduce the effective persistence relative to the monthly or quarterly frequencies typical of macroeconomic data. Nevertheless, with $T = 17$, the finite-sample component remains a concern. To address this, we apply the analytical bias correction proposed by Herbst and Johannsen (2024), who derive the leading-order bias of LPs impulse response estimators and provide plug-in corrections that eliminate bias of order $O(T^{-1})$.

The key insight is as follows. Herbst and Johannsen (2024) derive the bias under the assumption that shocks are exogenous (i.e., for LP-OLS). In our setting, we first recover exogenous variation in fishing revenue through the shift-share IV described in Section 4.2.1, and then apply the bias correction to the resulting LP-IV estimates. In a panel LPs with fixed effects estimated over N entities and T time periods, the bias of the estimator $\hat{\beta}^h_{LS,N}$ at horizon h takes the following form when lagged outcome controls are included:[5]

$$E[\hat{\beta}^h_{LS,N}] - \beta^h = -\frac{1}{T-h}\sum_{j=1}^{h}[1+\varpi_j]\,\beta^{h-j} + O(T^{-3/2}), \tag{4.4}$$

where ϖ_j captures cross-entity covariance terms that attenuate rapidly as N grows. In our setting with large N, the ϖ_j terms become negligible, and the bias simplifies to:

$$-\frac{1}{T-h}\sum_{j=1}^{h}\beta^{h-j}. \tag{4.5}$$

This remaining bias is driven entirely by the estimation of community fixed effects with finite T, analogous to the Nickell bias in standard dynamic panel models. Applying the recursive correction of Herbst and Johannsen (2024) to our LP-IV estimates, we construct the bias-corrected estimator (LP-IV-BCC):

$$\hat{\beta}^h_{\text{BCC},N} \;=\; \hat{\beta}^h_{IV,N} \;+\; \frac{1}{T-h}\sum_{j=1}^{h}\hat{\beta}^{h-j}_{\text{BCC},N}, \qquad h \geq 1, \tag{4.6}$$

with the anchor $\hat{\beta}^0_{\text{BCC},N} = \hat{\beta}^0_{IV,N}$, where $\hat{\beta}^h_{IV,N}$ denotes the LP-IV estimate at horizon h. The correction works recursively: at each horizon h, the estimator adds back the estimated bias using previously corrected coefficients. We refer to this estimator as LP-IV-BCC throughout. A detailed derivation is provided in Appendix C. For inference, we combine lag-augmented controls with cluster-robust standard errors at the community level, following Olea and Plagborg-Møller (2021), Herbst and Johannsen (2024), and Jordà and Taylor (2025).

4.2.3 Monte Carlo Simulation

To validate whether our approach recovers impulse response functions that accurately track cumulative dynamic multiplier effects, we design Monte Carlo simulations based on a structural VAR data-generating

[5] Herbst and Johannsen (2021, 2024) also derive the bias without lag controls, showing that in this case the bias depends on IRF coefficients across all horizons, not only past ones.

process calibrated to resemble the relationship between local economic outcomes and fishing performance in our data, matching the finite-sample constraints of small T and large N. The simulations focus on four questions: (1) whether including lagged outcomes as controls affects IRF bias, as predicted by Herbst and Johannsen (2021, 2024); (2) the comparative performance of level versus long-difference specifications as demonstrated by Piger and Stockwell (2025); (3) how the persistence of the outcome variable governs the severity of the bias; and (4) how well the bias correction of Herbst and Johannsen (2021, 2024) recovers the true IRFs relative to uncorrected estimates.

The simulations confirm our preferred specification: long-difference outcome, inclusion of lagged outcome and treatment variables as controls, and recursive bias correction. We also select the lag lengths $p = q = 2$ based on these simulations, balancing the bias-reduction benefits of lag augmentation against the cost of consuming limited time-dimensional observations.[6] This combination substantially reduces finite-sample bias while maintaining adequate precision. In particular, we demonstrate that inclusion of the lagged dependent variable is critical for identification, a feature absent from Watson et al. (2021). Our simulations show that even when the contemporaneous shock is exogenous, the estimated IRF is biased without these controls, consistent with the incidental parameter problem discussed in Herbst and Johannsen (2024) and Klosin (2024). This demonstrates that, although we adopt the identification strategy of Watson et al. (2021), we improve upon their estimation not only for long-run dynamic multiplier effects but also for the contemporaneous effects they originally estimated. Detailed simulation design and results are provided in Appendix E.

4.2.4 Data

We compile a dataset of economic and commercial fishing variables for Alaskan communities that engaged in commercial fisheries over the period 2000–2016, following Watson et al. (2021).

Data on permit-owner earnings are sourced from the Alaska Commercial Fisheries Entry Commission (CFEC) Basic Information Tables, which provide annual records of harvests and earnings for each community-fishery pair. Alaskan fisheries are defined by species, fishing district, and gear type, and participation requires a fishery-specific permit issued by the CFEC. In 2010, 20,275 permits were issued across

[6] Our simulations indicate that even a single lag ($p = q = 1$) substantially reduces bias. We adopt two lags for additional robustness, given that the modest loss of degrees of freedom remains acceptable. Standard information criteria are not well suited for lag selection in this context: AIC and BIC are invalid for LPs regressions at $h > 1$ due to the moving average structure of the residuals (Jordà and Taylor, 2025), and in our finite-sample environment they select zero lags ($p = q = 0$). We therefore rely on our Monte Carlo simulations to determine that two lags are sufficient.

205 fisheries, with permit-owner communities identified by the addresses listed on their permits. The value of deliveries to local processors is aggregated from individual delivery records reported through the Alaska Department of Fish and Game's (ADF&G) fish tickets and eLandings systems.

Information on local wages and employment comes from the Alaska Local and Regional Information (ALARI) database, which provides sector-specific economic data for 344 Alaskan communities. Unlike the Bureau of Labor Statistics' Quarterly Census of Employment and Wages (QCEW), which reports statistics by place of work at the county level, ALARI reports wage and employment data by workers' places of residence at the community level. This is enabled by the Alaska Department of Labor and Workforce Development's integration of unemployment insurance records with data from Alaska's Permanent Fund Dividend application database, which includes applicants' residential addresses (Watson et al., 2021). This residence-based, community-level measurement is essential for our analysis, as county-level statistics would obscure meaningful variation across heterogeneous communities (Guettabi and James, 2020; Jacobsen et al., 2023).

Two features of the ALARI data warrant note. First, because records are derived from unemployment insurance data, they exclude commercial fishers and crew (who are typically self-employed or contract workers), as well as other self-employed individuals. Wage and employment data for workers employed by commercial processors are included. This means our outcome variables capture spillover effects into non-fishing sectors rather than direct effects on fishing employment. Second, ALARI measures the number of individual workers rather than full-time equivalent (FTE) jobs. Given the highly seasonal nature of Alaska's seafood industry, our employment estimates may differ from studies that rely on FTE measures.

Because our specification uses log differences of fishing revenue as the shock variable and long differences of outcomes, each observation requires at least two consecutive years of non-missing data. The inclusion of two lags of both the outcome and shock variables further reduces the usable sample. Communities with intermittent or minimal fishing activity are therefore excluded from the estimation sample. As a result, the effective sample for resident earnings consists of approximately 150 communities at $h = 0$, declining to approximately 130 at $h = 5$ as the forward-looking long difference requires additional lead observations. These communities represent those most consistently engaged in commercial fisheries over the sample period. Summary statistics for the estimation sample are reported in Appendix Table G.1.

4.3 Results

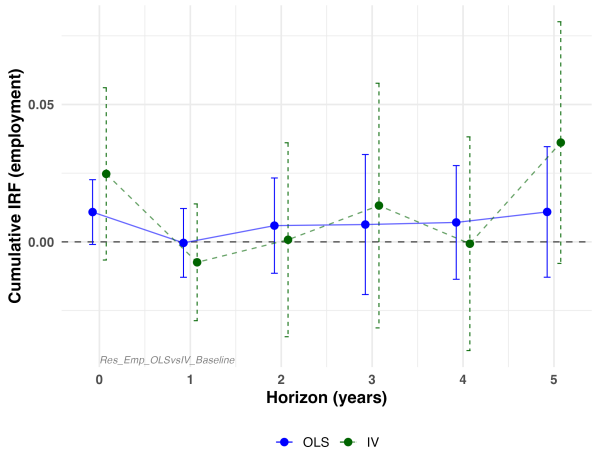
We present the estimated impulse response functions from our panel local projection framework, capturing the dynamic response of non-fishing employment and wages to shocks in fishing revenue over horizons $h = 0, \ldots, 5$. We report estimates under three specifications: LP-OLS, LP-IV, and LP-IV-BCC, each presented separately for resident earnings and local landings measures of fishing activity.

4.3.1 Spillover Effects: Resident Earnings versus Local Landings

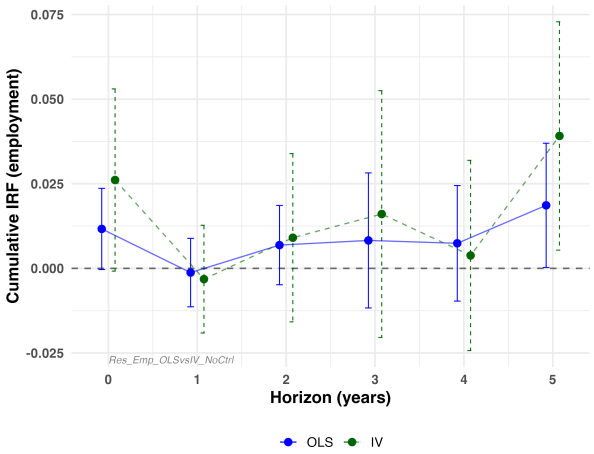
Following Watson et al. (2021), we separately estimate effects from two channels through which fishing activity enters a community: resident permit-owner earnings, which measure fishing revenue attributed to the community where the permit holder resides, and local landings, which measure the value of fish delivered to processors located in the community. These two channels may generate different spillover dynamics, as they involve different actors and income flows within the local economy.

Resident earnings

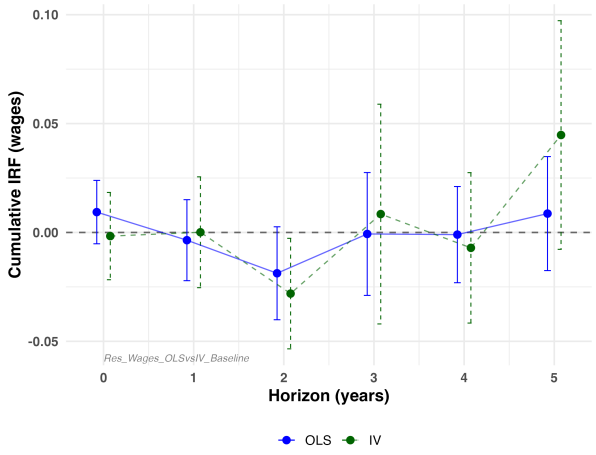
Figure 4.1 compares our preferred baseline specification, a dynamic panel local projection that includes lagged outcome and treatment controls (as motivated by our Monte Carlo simulations in Appendix E and the theoretical findings of Klosin (2024) and Herbst and Johannsen (2024)), against a no-control specification that omits these lags. Both are estimated as local projections over horizons $h = 0, \ldots, 5$; they differ only in whether lagged outcome and treatment terms are included. Although both specifications exhibit broadly similar impulse response paths, they differ in magnitude. These results indicate that including lagged outcome controls can meaningfully shift the estimated response, consistent with the bias mechanism in dynamic panels characterized by Klosin (2024) and Herbst and Johannsen (2024). Furthermore, LP-OLS and LP-IV estimates differ meaningfully across horizons, with OLS estimates exhibiting bias from the endogeneity of community-level fishing performance that the shift-share instrument addresses.



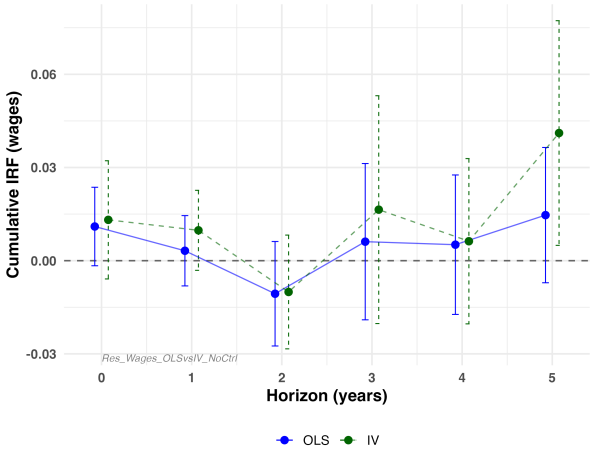
(a) Employment, Baseline



(b) Employment, No Controls



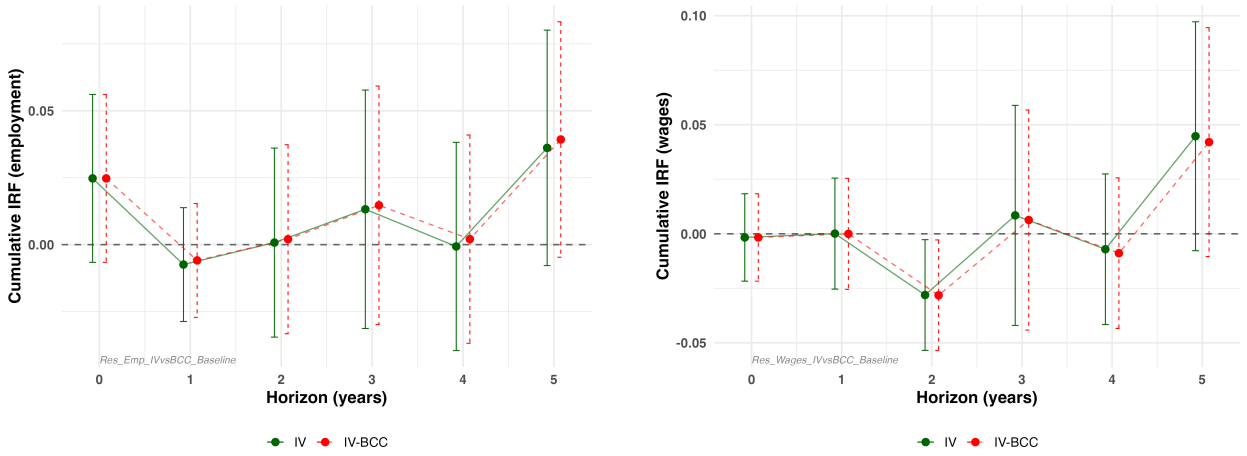
(c) Wages, Baseline



(d) Wages, No Controls

Figure 4.1: LP-OLS versus LP-IV impulse response of wages and employment to resident earnings. Left panels show the baseline specification; right panels show the no-control specification. Blue lines denote LP-OLS; green lines denote LP-IV. Shaded areas represent 90% confidence intervals.

Figure 4.2 compares the LP-IV and LP-IV-BCC estimates of cumulative wage and employment growth. The bias correction produces only modest adjustments, suggesting that the long-difference specification already substantially mitigates persistence-driven bias, or that the autoregressive persistence of local economic outcomes is not sufficiently high to generate substantial bias in this specification.



(a) Employment

(b) Wages

Figure 4.2: LP-IV versus bias-corrected (LP-IV-BCC) impulse response of wages and employment to resident earnings (baseline specification, first-stage F-statistic = 90.12). Green solid line denotes LP-IV; red dashed line denotes LP-IV-BCC following Herbst and Johannsen (2024). Shaded areas represent 90% confidence intervals.

For total employment, we find a positive effect at impact ($h = 0$) near the 10% significance threshold, followed by statistically insignificant effects over intermediate horizons ($h = 1$–$h = 4$) and a positive effect at the maximum horizon ($h = 5$). Wages exhibit a similar non-monotonic pattern but with no effect at impact and a significant negative response at $h = 2$, possibly reflecting compositional shifts in the employed workforce or differential short-run adjustment between labor quantities and compensation (Figure 4.2b). At $h = 5$, wages also turn positive, paralleling the employment results.

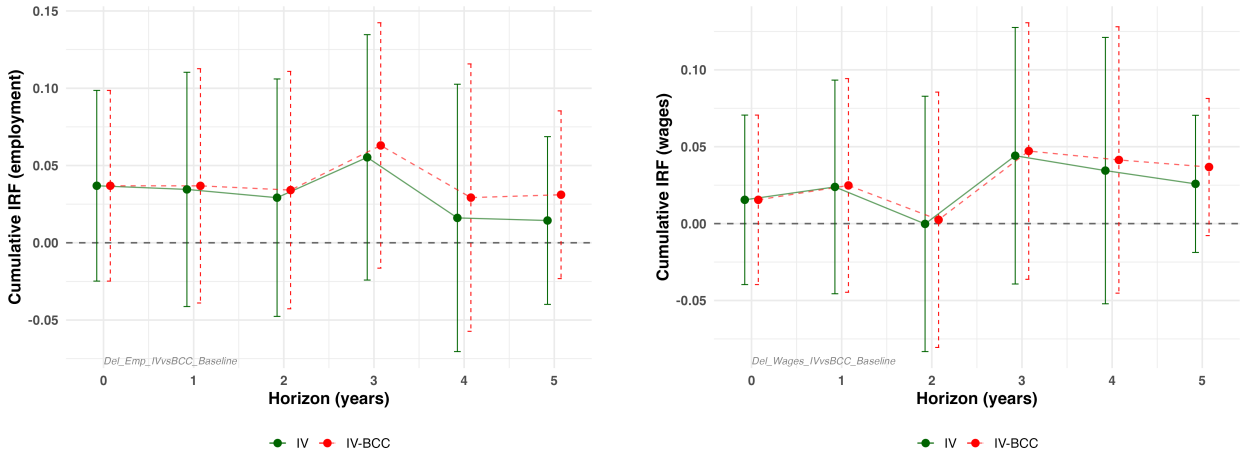
However, two caveats constrain the interpretation of the $h = 5$ estimates. First, our panel's time dimension ($T = 17$) precludes estimation beyond $h = 5$ due to finite sample limitations, so we cannot determine whether this point reflects a stabilized response, a transitory peak, or an ongoing trajectory. Second, as the terminal horizon, $h = 5$ is estimated on the smallest effective sample and is most vulnerable to endpoint sensitivity, with wider confidence intervals relative to other horizons. We therefore interpret $h = 5$ not as a steady-state long-run multiplier, but as the cumulative response at the maximum observable horizon.

The bias correction of Herbst and Johannsen (2024) eliminates bias of order $O(T^{-1})$ analytically but does not remove higher-order terms, and a residual negative bias generally persists (Herbst and Johannsen, 2024). Our Monte Carlo simulations calibrated to our data setup (Appendix E) confirm that a mild residual downward bias, approximately 15% of the true IRF at longer horizons, remains after correction, consistent with the $O(T^{-3/2})$ remainder in Equation (4.4). Because this residual bias is signed downward rather than

mean-zero, the corrected estimates in Figure 4.2 represent conservative lower bounds on the true cumulative multipliers. Taken together, the evidence is consistent with, though does not definitively establish, non-negative cumulative spillovers from resident earnings to local employment and wages over the observable horizon.

Local landings

We also estimate dynamic multiplier effects using local landings revenue. Figure 4.3 shows that the estimated effects on both employment and wage growth are generally positive but broadly insignificant across all horizons. The muted spillovers are consistent with the finding of Watson et al. (2021) that processing labor at landing sites is predominantly composed of non-resident seasonal workers whose earnings leak out of the local economy rather than generating induced effects. The bias correction reveals a modest but increasing downward bias in the uncorrected IV estimates at longer horizons for both employment and wages, consistent with the theoretical predictions of Herbst and Johannsen (2024) and our MC simulation results.



(a) Employment **(b)** Wages

Figure 4.3: LP-IV versus bias-corrected (LP-IV-BCC) impulse response of wages and employment to local landings revenue (baseline specification, first-stage F-statistic = 37.82). Green solid line denotes LP-IV; red dashed line denotes LP-IV-BCC following Herbst and Johannsen (2024). Shaded areas represent 90% confidence intervals.

Taken together, the local landings estimates are generally positive but statistically indistinguishable from zero at most individual horizons, consistent with weaker spillover potential from this channel relative to resident earnings at the aggregate employment and wage level. In the next section, we decompose these aggregate effects into traded and non-traded sectors to examine whether the two channels differ in their

sectoral composition.

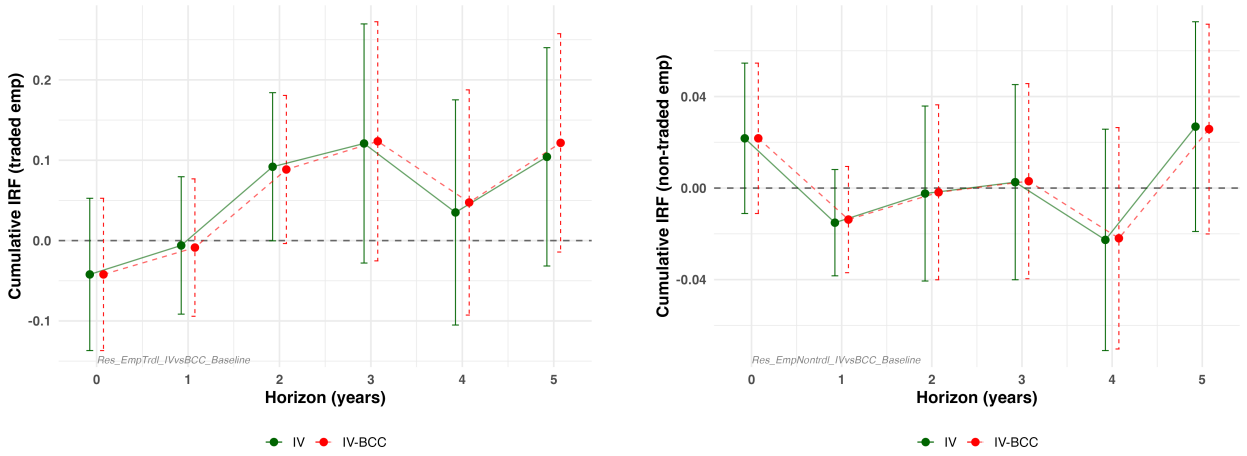
4.3.2 Sectoral Decomposition: Traded and Non-traded Employment

We decompose the aggregate employment response into traded and non-traded sectors to examine the channels through which fishing revenue spillovers operate. Following Moretti (2010) and Watson et al. (2021), we classify ALARI's industry-level employment into traded sectors (agriculture, natural resources, mining, and manufacturing, including fish processing) and non-traded sectors (all other industries). Because the two sectors may respond to fishing revenue shocks with different timing and magnitude, the dynamic framework allows us to examine their relative contributions across horizons.

Sectoral decomposition: Resident earnings

Figure 4.4 decomposes the resident earnings employment response into traded and non-traded sectors. At impact ($h = 0$), the traded sector response is near zero, but it builds over subsequent horizons, approaching significance at $h = 2$ and re-emerging at $h = 5$. In contrast, the non-traded sector response is statistically insignificant across all horizons. The contrasting trajectories suggest that resident earnings spillovers operate primarily through traded-sector channels rather than through induced spending in non-traded sectors. We discuss this point further, including the overlap between permit-owner residence and local processing activity, in Section 4.4.

We return to this pattern in Section 4.3.4, where the common-sample GMM specification reveals a jointly significant non-traded employment response, indicating that the sectoral composition and the magnitude of spillovers may be sample- and period-dependent (in particular, the common-sample restriction concentrates the analysis on communities consistently engaged in commercial fishing throughout the period). Possible interpretations of these patterns are discussed in Section 4.4.



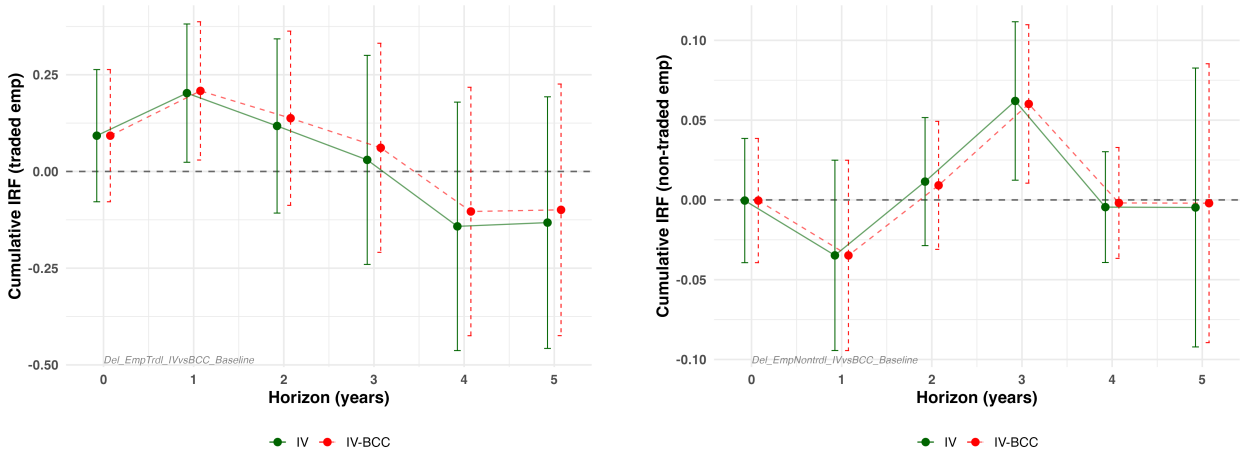
(a) Traded sector employment

(b) Non-traded sector employment

Figure 4.4: LP-IV versus bias-corrected (LP-IV-BCC) impulse response of traded and non-traded sector employment to resident earnings (baseline specification, first-stage F = 115.1 and 98.81, respectively). Green solid line denotes LP-IV; red dashed line denotes LP-IV-BCC following Herbst and Johannsen (2024). Shaded areas represent 90% confidence intervals.

Sectoral decomposition: Local landings

Figure 4.5 decomposes the local landings employment response into traded and non-traded sectors. The traded sector shows a short-run positive effect that reaches statistical significance at $h = 1$, then quickly diminishes and becomes statistically indistinguishable from zero at longer horizons. The non-traded sector exhibits a down-up-swing trajectory: a brief negative dip at $h = 1$, a positive peak at $h = 3$ that reaches pointwise significance, and insignificance at the remaining horizons. As we show in Section 4.3.4, this dynamic path is jointly significant at the 1% level despite only one individual horizon reaching conventional significance.



(a) Traded sector employment

(b) Non-traded sector employment

Figure 4.5: LP-IV versus bias-corrected (LP-IV-BCC) impulse response of traded and non-traded sector employment to local landings revenue (baseline specification, first-stage F = 31.15 and 37.39, respectively). Green solid line denotes LP-IV; red dashed line denotes LP-IV-BCC following Herbst and Johannsen (2024). Shaded areas represent 90% confidence intervals.

The local landings sectoral pattern differs notably from resident earnings. Traded sector employment responds relatively rapidly and then fades, rather than building gradually over multiple years; non-traded employment exhibits a pronounced swing rather than the flat response observed for resident earnings. These contrasting dynamics, together with the muted aggregate effects in Figure 4.3, suggest that local landings shocks generate more concentrated short-run activity rather than sustained spillovers. The magnitude of these effects may also be understated, since ALARI excludes seasonal and contract labor (Section 4.2.4), which constitutes a substantial share of the seafood processing workforce. We discuss possible interpretations of the resident earnings versus local landings contrast and its implications in Section 4.4.

4.3.3 Heterogeneity by Fishing Dependency

We examine whether the dynamic multiplier effects differ across communities with varying levels of dependence on commercial fisheries. Following Watson et al. (2021), we construct a fishing dependency ratio defined as the community-level mean of resident fishing revenue divided by total wage income, and progressively restrict the estimation sample to communities above the 50th, 60th, 70th, 80th, and 90th percentiles of this distribution.

Figure 4.6 presents the results for resident earnings wages and employment. Across all dependency thresholds, the impulse response paths follow broadly similar trajectories for both outcomes, preserving the

same qualitative pattern observed in the full sample (Figure 4.2). However, highly fishing-dependent communities exhibit notably larger amplitude oscillations at intermediate horizons, suggesting greater sensitivity to fishing revenue shocks during the adjustment process. This amplification is intuitive: in communities where fishing revenue accounts for a larger share of the local economy, the same percentage shock translates into a proportionally larger impact on total economic activity, predicting larger swings in both directions even when the underlying spillover mechanism is similar. By $h = 5$, all subsamples converge toward positive effects regardless of dependency level. The widening of confidence intervals at higher thresholds reflects the reduction in sample size rather than a qualitative change in the estimates.

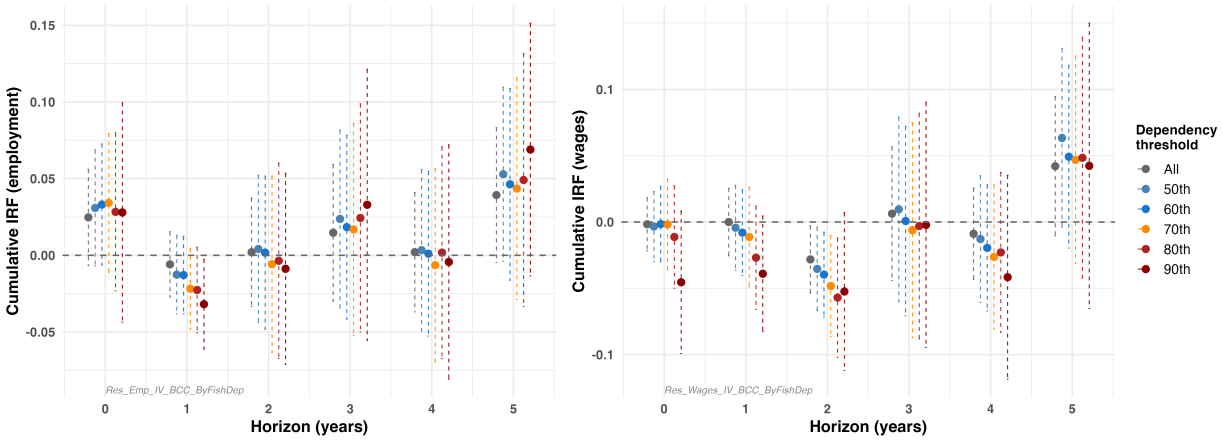


Figure 4.6: LP-IV-BCC impulse response of wages and employment to resident earnings by fishing dependency. Subsamples progressively restrict to communities above the indicated percentile threshold of the dependency distribution (community-level mean ratio of fishing revenue to total wages). Dots represent point estimates; vertical bars represent 90% confidence intervals.

A notable finding from this decomposition is the emergence of clearer negative employment and wage effects at short horizons among highly fishing-dependent subsamples, effects that were not apparent in the pooled estimation using the full sample. Even at $h = 0$ and $h = 1$, where the pooled-sample estimate is near zero, high-dependency subsamples display visibly negative responses despite the reduced sample size from subgrouping. This suggests that fishing revenue shocks may not operate uniformly in a positive direction in the short run, particularly for communities most exposed to the fishing sector. A possible interpretation is short-run labor crowding-out: during revenue surges, workers temporarily reallocate from formal non-fishing employment toward fishing-related activities, reducing measured employment and wage income in non-fishing sectors. Because ALARI excludes self-employed fishing earnings, this reallocation registers as a decline in the measured outcomes even though total household income, including fishing earnings, may

be at least constant or rising. This effect is plausibly stronger in fishing-dependent communities, where the fishing sector commands a larger share of the local labor pool. Consistent with this interpretation, the non-traded sector employment response in Figure 4.4b exhibits a negative dip at the same short horizon ($h = 1$), which may indicate that the crowding-out operates primarily through non-traded sectors where labor is most readily reallocated toward fishing activity. A similar negative dip appears in the local landings non-traded employment response (Figure 4.5b), suggesting that this short-run reallocation pattern is not specific to the resident earnings channel.

Overall, higher fishing dependency amplifies the *volatility* of the dynamic response but does not systematically alter its *direction*: more dependent communities experience larger swings during adjustment, including a more visible short-horizon negative effect, while the long-run spillover trajectory operates similarly across dependency levels.

4.3.4 Joint Inference on the Impulse Response Path

The pointwise estimates presented in the preceding sections trace the dynamic trajectory of fishing revenue shocks at each horizon, but they have an inherent limitation: because each horizon is estimated as a separate regression, the cross-horizon covariance structure is not recovered, and pointwise confidence intervals reflect only the marginal uncertainty at each horizon in isolation. However, estimates across horizons are correlated, since they are derived from overlapping samples of the same communities and time periods, and this correlation structure contains information that pointwise inference discards. This makes it difficult to assess whether the impulse response path as a whole constitutes statistically significant evidence of a dynamic effect.[7]

To address this, Jordà and Taylor (2025) cast the LP system as a GMM problem, stacking the moment conditions across all horizons to obtain a joint covariance matrix that enables testing whether the entire impulse response path is jointly significant. Following their approach, we test the null hypothesis $H_0 : \beta^0 = \beta^1 = \cdots = \beta^5 = 0$, which asks whether the full dynamic response to a fishing revenue shock is jointly distinguishable from zero, even when individual horizons are not. Because the system GMM requires a common sample in which all horizon outcomes are simultaneously observed, the effective sample size is

[7] Jordà and Taylor (2025) draw an analogy to multicollinearity in OLS regression: just as individual t-statistics on correlated regressors can be insignificant while the F-test on the overall model overwhelmingly rejects the joint null, pointwise bands on an impulse response can be wide at every horizon even when the entire path is jointly significant. Relying solely on horizon-by-horizon inference may therefore understate the overall statistical evidence for a dynamic effect.

reduced by up to approximately 40% at shorter horizons relative to the horizon-specific regressions, yielding less precise point estimates with wider confidence intervals. Nevertheless, this approach serves two purposes: first, it provides an independent validation of the multiplier trajectories found in the main results under a different sample composition; and second, it enables joint inference on whether the overall dynamic effects are statistically meaningful. Figures F.1 and F.2 in Appendix F show that the GMM and individual LP-IV estimates exhibit broadly similar qualitative trajectories and converge at $h = 5$ where the two samples are nearly identical, supporting the use of the GMM approach for joint inference despite the reduced sample.

Table 4.1 reports the results. For resident earnings, total employment is jointly significant at the 10% level, and non-traded employment is significant at the 5% level, despite no individual horizon reaching conventional significance for either outcome. For total employment, whose pointwise estimates exhibit a generally upward trajectory toward positive effects (Figure 4.2a), the joint test provides complementary evidence that fishing revenue shocks generate a meaningful dynamic effect over the full horizon, even though horizon-specific inference remains statistically ambiguous at each individual point. We also find joint significance for resident earnings non-traded employment, an outcome for which the full-sample pointwise estimates in Figure 4.4b showed no significance.

$H_0 : \boldsymbol{\beta} = \mathbf{0}$	$\chi^2(6)$	p-value
Panel A: Resident earnings		
Wages	5.13	0.527
Employment	11.19	0.083*
Employment (traded)	5.54	0.477
Employment (non-traded)	14.36	0.026**
Panel B: Local landings		
Wages	2.96	0.814
Employment	4.36	0.629
Employment (traded)	4.06	0.669
Employment (non-traded)	17.22	0.009***

***, **, * denote significance at the 1%, 5%, and 10% levels.
Test statistics are cluster-robust at the community level.
Estimation uses the common sample in which all horizons $h = 0, \ldots, 5$ are simultaneously non-missing.

Table 4.1: Joint Wald test for the full impulse response path ($h = 0, \ldots, 5$) under the common-sample GMM specification (effective period: 2006–2016).

Closer inspection of the GMM impulse responses (Appendix F, Figure F.1) suggests that the joint significance for both total and non-traded employment is driven primarily by visibly larger contemporaneous responses ($h = 0$) under the common-sample specification rather than by accumulated effects across longer

horizons. The contemporaneous non-traded spillover pattern echoes the finding of Watson et al. (2021), who report strong contemporaneous multiplier effects induced in non-traded sectors by resident earnings. Notably, our joint estimation also finds significant non-traded employment effects for local landings, extending this result beyond the resident earnings channel.

We further note an interpretive point about the GMM specification. Because the common-sample restriction simultaneously concentrates the analysis on a narrower time window (effectively 2006–2016) and on communities with consistent fishing activity across all horizons, the GMM estimates capture a different economic environment than the full-sample LP-IV estimates. The divergence between the two reflects two distinct sources: the common-sample restriction reduces the effective sample size and widens confidence intervals, lowering statistical power; and the same restriction changes the underlying economic environment captured by the estimates, both in time and in community composition. For wages and traded sector employment, the joint test finds no significant evidence of an overall dynamic effect, even though the pointwise traded employment estimates approached significance at $h = 2$ and $h = 3$ in the main results. However, the full-sample LP-IV estimates, which exploit the complete time span and all available communities at each horizon, provide a more representative picture of the dynamic response path. The joint GMM results should therefore be treated as supplementary evidence: they support the presence of dynamic effects where significance is found, but their null findings should be interpreted as inconclusive rather than as evidence against a dynamic effect.

Overall, the joint inference results provide complementary evidence for the main pointwise findings. For resident earnings, the joint test supports significant dynamic effects on total employment, broadly reinforcing the upward trajectory observed in the pointwise results. For local landings, the joint test identifies a significant dynamic effect on non-traded sector employment, consistent with the delayed positive response visible in the pointwise analysis. Wages and traded sector employment do not reach joint significance under either specification, which we interpret cautiously given the reduced power and the period concentration of the common-sample estimation discussed above. Finally, we note that the joint test uses the IV estimates without the bias correction of Herbst and Johannsen (2024); since the bias-corrected and uncorrected estimates are broadly similar in our main results, this is unlikely to qualitatively alter the conclusions.

4.3.5 Scaling Elasticities to Level Effects

The estimated coefficients $\hat{\beta}^h$ are elasticities representing cumulative percent changes. To provide a more concrete interpretation, we convert them into level changes following Watson et al. (2021): $\Delta Y/\$ = \hat{\beta}^h \cdot \bar{y}/\bar{x}$, where $\bar{y}$ and $\bar{x}$ are the sample means of the outcome and fishing revenue. For employment, this yields jobs per million dollars of fishing revenue; for wages, dollars of wage income per dollar of fishing revenue.

The full-scaled paths are reported in Appendix H. For resident earnings, the scaled total employment point estimate ranges from approximately 23 jobs per million dollars at impact to 30 jobs at $h = 5$, with traded sector employment building from near zero to about 7 jobs at longer horizons. Non-traded employment accounts for the bulk of the total point estimate, though, as noted in Section 4.3.2, it is not statistically distinguishable from zero at any individual horizon. The scaled wage effect reaches 0.63 dollars per dollar at $h = 5$; because formal-sector wages exclude self-employed fishing income, this represents spillover into the non-fishing economy. For local landings, scaled total employment ranges from 12 to 25 jobs per million dollars across horizons, with traded employment peaking near 6 jobs at impact and non-traded employment reaching 22 jobs at $h = 3$. Both specifications reproduce the dynamic patterns documented above, expressed here in level units that facilitate substantive interpretation.

4.4 Discussion

This paper provides the first estimates of dynamic multiplier effects from commercial fisheries on local economies, extending the contemporaneous causal evidence of Watson et al. (2021) and the time-series analysis of Seung (2008) to a multiyear horizon with causal identification at the community level. Using a panel local projections framework with bias correction for up to 160 Alaska fishing communities over 2000–2016, we trace the cumulative impulse response of local employment and wages to exogenous changes in fishing revenue over horizons up to five years.

Several broad patterns emerge from the estimation. First, the impulse response paths suggest that fishing revenue shocks generate effects that extend beyond the contemporaneous impact, and that these longer-horizon effects are non-monotonic rather than smoothly building or decaying, with short-term negative dips, particularly in wages, followed by positive effects at longer horizons. Furthermore, these dynamic effects are more clearly detected in employment than in wages: the full-sample pointwise estimates trace a generally

upward trajectory for employment with close to statistical significance at the maximum observable horizon, and the joint horizon test provides further support by rejecting the null of zero dynamic effects for total employment.

Second, the sectoral composition of spillovers differs across the two channels. Both resident earnings and local landings generate impacts on traded-sector employment, but with different temporal dynamics: resident earnings produce a delayed but gradually building positive response, whereas local landings generate a more immediate but transient effect that fades at longer horizons. This is somewhat surprising, as one might expect traded-sector effects to be driven primarily by local landings, where processing firms are located. That the more persistent traded-sector response instead arises from resident earnings suggests a channel we do not fully understand from our reduced-form estimates alone, though one possibility is that permit-owner residence and processor employment catchments overlap geographically in Alaska communities. For non-traded employment, local landings show a lagged positive response at intermediate horizons, consistent with local spending by processing workers, including seasonal workers, gradually translating into retail and service employment in nearby communities. However, this effect builds with a lag and dissipates before the maximum observable horizon, suggesting that the induced spending channel from local landings generates a temporary rather than sustained employment response in non-traded sectors.

Finally, higher fishing dependency amplifies the volatility of the dynamic response, including more pronounced short-term negative effects, without systematically altering its long-run direction. The short-term negative effects on non-fishing employment and wages, particularly visible in the non-traded sector and concentrated among highly fishing-dependent communities, are consistent with labor crowding-out: when fishing revenue increases, the harvesting sector absorbs additional labor from the local workforce, drawing workers away from formal non-fishing employment. Moretti (2010) shows that the magnitude of local employment multipliers depends critically on the elasticity of local labor supply, with less elastic labor markets producing stronger crowding-out when a dominant sector expands. The highly fishing-dependent communities in our sample exemplify this dynamic: they tend to be small, rural, and geographically remote, with highly constrained labor pools.[8] In such settings, external labor cannot readily flow in to absorb the increased demand from fishing, so the crowding-out of non-fishing sectors is more pronounced in the short run. As the labor market adjusts over subsequent years, whether through seasonal workforce turnover,

[8]Communities above the 90th percentile of fishing dependency have a median population of approximately 95 and median employment of 45 workers.

gradual in-migration, or induced local demand, the initial negative effect gives way to the positive spillovers observed at longer horizons. The long-run trajectories across dependency levels all shift toward positive effects at the maximum observable horizon, suggesting that this adjustment ultimately occurs regardless of initial fishing dependence, but the transition is more disruptive where the local labor pool is smaller and fewer non-fishing sectors are available to buffer the shock.

These findings carry several implications for fisheries management and coastal economic policy. Most fundamentally, the evidence that fishing revenue shocks generate employment effects extending beyond the contemporaneous period suggests that the economic returns to fishery investments, such as stock rebuilding, habitat restoration, and conservation programs, may extend well beyond the harvesting sector and beyond the short run. Because fisheries are a renewable resource, successful stock recovery can sustain improved productivity on a continuing basis, meaning that the local economic spillovers documented here are not one-time gains but potentially recurring benefits under sustainable management.

The different temporal dynamics across channels and sectors have practical consequences for policy design. Our results suggest that policymakers should consider not only the magnitude but also the timing and sectoral composition of expected employment effects when evaluating fisheries interventions. For instance, a stock rebuilding program that enhances landings and strengthens resident earnings may show little immediate employment response but generate gradually building traded-sector gains over multiple years, whereas policies that increase local landings volume may produce quicker but more transient employment effects concentrated in processing. Evaluating either type of intervention after only one or two years risks premature conclusions, particularly given our finding that the non-monotonic trajectory implies short-term dips in wages and employment, especially in highly fishing-dependent communities, may precede longer-run positive effects. Policy evaluation frameworks should therefore adopt time horizons consistent with the dynamic structure documented here, rather than relying on short-window assessments that risk understating the dynamic cumulative outcomes in local fishing economies to conservation and rebuilding investments.

Moreover, the distinct temporal dynamics across channels and sectors suggest that the effectiveness of fisheries policy depends on which spillover pathway it activates. Policies that strengthen resident permit-owner earnings, such as community quota programs or measures that encourage local ownership, may generate more durable employment gains through the traded sector, albeit with a delayed onset. By contrast, investments in local processing infrastructure can produce more immediate traded-sector employment through the local landings channel, though these effects appear transient; however, the same channel also

generates lagged positive effects on non-traded employment, such as retail and services, suggesting that processing activity eventually spills over into the broader local economy. Policymakers designing community economic development strategies around fisheries could therefore benefit from combining both channels: supporting resident ownership for sustained long-term horizon spillovers while investing in processing capacity for shorter-term employment gains and eventual non-traded sector growth. Finally, the amplified short-term volatility observed in highly fishing-dependent communities suggests that transitional support may be warranted during adjustment periods following major policy changes, even when the long-run economic trajectory is favorable.

Finally, the finding that highly fishing-dependent communities experience negative effects at early horizons underscores the importance of understanding community-specific context before implementing fisheries policy changes. Although the long-run trajectories across dependency levels all shift toward positive effects at the maximum observable horizon, the adjustment path is considerably more disruptive where fishing accounts for a larger share of the local economy. In small, remote communities with constrained labor pools, a fishing boom may induce short-run crowding-out in non-fishing sectors, as our heterogeneity analysis suggests. These patterns point to the potential value of assessing community-specific industrial structure and labor market conditions prior to major policy changes, and of considering targeted, time-limited adjustment assistance for highly dependent communities during the transition period.

Several limitations, however, should be noted. First, our reduced-form framework estimates the dynamic trajectory of multiplier effects but does not directly identify the mechanisms generating the observed patterns. The interpretations we offer, including labor crowding-out, the geographic overlap between permit-owner residence and processor catchments, and sectoral adjustment frictions, are plausible accounts consistent with our estimates rather than mechanisms we directly test. Investigating them through structural models that incorporate labor market frictions, seasonal employment dynamics, and inter-industry input-output linkages is an important direction for future research. Second, our panel covers 17 years (2000–2016), which limits the maximum estimation horizon to five years. We cannot determine whether the positive effects observed at the maximum horizon represent a stabilized long-run response, a transitory peak, or part of a continuing trajectory. Third, as the terminal horizon in a finite sample, the maximum horizon is estimated on the smallest effective sample and is most vulnerable to endpoint sensitivity, though the consistency of this pattern across fishing dependency subsamples provides some reassurance. Finally, while we apply analytical bias correction and validate our approach through Monte Carlo simulations, a mild residual downward bias

persists after correction, so the reported estimates should be interpreted as conservative lower bounds on the true cumulative multiplier effects.

Appendices

Appendix

A Incidental Parameter Problem (IPP) and Instrumental Variable (IV) Approach

Suppose we have a simple panel model with individual fixed effects, where x_{it} is endogenous:

$$y_{it} = \alpha_i + \beta x_{it} + \epsilon_{it}, \quad i = 1, \ldots, N, \quad t = 1, \ldots, T.$$

We want to estimate β using an instrumental variable z_{it} for x_{it}. In an ideal IV framework (if α_i were known), we would have the moment condition:

$$E\Big[z_{it}\Big(y_{it} - \alpha_i - \beta x_{it}\Big)\Big] = 0.$$

However, in practice α_i is not observed. Let the estimated fixed effect be $\hat{\alpha}_i = \alpha_i + \eta_i$, where η_i is the estimation error of order $O(1/T)$ that does not vanish as $N \to \infty$ if T is fixed. Plugging $\hat{\alpha}_i$ into the moment condition yields:

$$\frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T} z_{it}\Big[\Big(y_{it} - \alpha_i - \beta x_{it}\Big) - \eta_i\Big] = 0.$$

The first term has expectation zero by IV validity. However, the second term introduces:

$$-\frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T} z_{it}\,\eta_i,$$

which does not vanish when T is small. Via Taylor expansion, the bias in $\hat{\beta}$ is:

$$\hat{\beta} - \beta_0 \approx \left\{\frac{\partial m}{\partial \beta}\Big|_{\beta_0}\right\}^{-1} \cdot \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T} z_{it}\,\eta_i.$$

Since η_i is of order $O(1/T)$ and does not vanish when T is fixed (Nickell, 1981), the bias in $\hat{\beta}$ is non-negligible. Even though IV methods successfully address endogeneity, they cannot remove the bias resulting from imprecise estimation of fixed effects. Therefore, even if, the treatment is exogenous, due to finitie-time sample, bias still exists as demonstrated in our MC simulation in Appendix E.

Assuming $N \to \infty$ with fixed T:

$$\frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T} z_{it}\,\eta_i \xrightarrow{p} E\left[\left(\frac{1}{T}\sum_{t=1}^{T} z_{it}\right)\eta_i\right] \neq 0.$$

This incidental parameter problem is not an issue in static linear panel data (as used in Watson et al. (2021)) since the nuisance parameters α_i are differenced out by within or first-difference transformation. However, once the panel involves dynamic structure (e.g., autogressive DGP), the IPP arises even after using IV.

B Nickell Bias in Dynamic Panels under Finite T and Large N

Consider the dynamic panel model:

$$y_{it} = \rho y_{i,t-1} + \alpha_i + \epsilon_{it}, \quad i = 1, \ldots, N, \quad t = 1, \ldots, T.$$

The within transformation eliminates α_i but creates:

$$\tilde{y}_{i,t-1} = \left(1 - \frac{1}{T}\right) y_{i,t-1} - \frac{1}{T}\sum_{s \neq t} y_{i,s-1}, \quad \tilde{\epsilon}_{it} = \left(1 - \frac{1}{T}\right)\epsilon_{it} - \frac{1}{T}\sum_{s \neq t}\epsilon_{is}.$$

Because the same observations form both individual values and averages, the overlap induces:

$$E\left[\tilde{y}_{i,t-1}\,\tilde{\epsilon}_{it}\right] \neq 0,$$

which downward-biases the estimation of ρ. Nickell (1981) showed:

$$E(\hat{\rho}) - \rho \approx -\frac{1+\rho}{T-1}.$$

The bias is of order $O(1/T)$ and remains even as N increases. In our LPs specification with log-long-difference in outcome variable (Equation 4.2), the outcome generally has reduced persistence after differencing, mitigating the severity of the Nickell bias. However, the bias from fixed T, even if moderate, remains a concern. As shown in Appendix A, IV does not automatically recover consistency.

C Detailed Derivation of Approximate Bias Correction

This appendix provides the detailed analytical derivation of the bias correction applied in the main text, following Herbst and Johannsen (2021, 2024). We begin with the time-series case and then extend to the panel setting. Throughout this appendix, β^h denotes the impulse response coefficient at horizon h, consistent with the main text.

Time-Series Local Projections. For a scalar outcome, the h-horizon LPs regression is:

$$y_{t+h} = \alpha^h + \phi_h' x_t + u_t^h, \qquad x_t = \left[\varepsilon_t,\ c_{t-1}'\right]', \quad h = 0, 1, \ldots,$$

where ε_t is the structural shock, c_{t-1} contains lagged controls, and ϕ_h is the full coefficient vector on x_t.[9] The inclusion of the lagged outcome y_{t-1} in c_{t-1} plays a significant role in controlling bias, as discussed in Appendix D later. The parameter of interest is $\beta^h \equiv \phi_{h,1}$, the first element of ϕ_h corresponding to ε_t, and $\hat{\beta}^h_{\mathrm{LS}}$ denotes its OLS estimator. The leading-order bias is:

$$E[\hat{\beta}^h_{\mathrm{LS}}] - \beta^h = \frac{1}{T-h} \sum_{j=1}^{h} \left[1 + \mathrm{tr}\{\Sigma_{c,0}^{-1}\Sigma_{c,j}\}\right]\beta^{h-j} + O(T^{-3/2}), \tag{C.1}$$

where $\Sigma_{c,j} = E[(c_{t-j} - \bar{c})(c_t - \bar{c})']$. When α^h is known:

$$E[\hat{\beta}^h_{\mathrm{LS}}] - \beta^h = -\frac{1}{T-h} \sum_{j=1}^{h} \mathrm{tr}\{\Sigma_{c,0}^{-1}\Sigma_{c,j}\}\beta^{h-j} + O(T^{-3/2}). \tag{C.2}$$

Herbst and Johannsen (2024) propose an approximate bias-corrected estimator:

$$\hat{\beta}^h_{\mathrm{BC}} = \hat{\beta}^h_{\mathrm{LS}} + \frac{1}{T-h} \sum_{j=1}^{h} \left[1 + \mathrm{tr}\left\{\widehat{\Sigma}_{c,0}^{-1}\widehat{\Sigma}_{c,j}\right\}\right] \hat{\beta}^{h-j}_{\mathrm{LS}}, \tag{C.3}$$

and a recursive refinement:

$$\hat{\beta}^h_{\mathrm{BCC}} = \hat{\beta}^h_{\mathrm{LS}} + \frac{1}{T-h} \sum_{j=1}^{h} \left[1 + \mathrm{tr}\{\widehat{\Sigma}_{c,0}^{-1}\widehat{\Sigma}_{c,j}\}\right] \hat{\beta}^{h-j}_{\mathrm{BCC}}, \qquad h \geq 1, \tag{C.4}$$

with $\hat{\beta}^0_{\mathrm{BCC}} = \hat{\beta}^0_{\mathrm{LS}}$.

[9] Herbst and Johannsen (2021, 2024) note that their results also apply when y_{t+h} is replaced with long differences $y_{t+h} - y_{t-1}$.

Extension to Panels with Fixed Effects. For N panel units $i = 1, \ldots, N$:

$$y_{i,t+h} = \mu_i^h + \beta^h \varepsilon_{i,t} + \tilde{\phi}_h' c_{i,t-1} + \gamma_h' c_{t-1} + u_{i,t}^h,$$

where μ_i^h is a horizon-specific individual fixed effect and $\tilde{\phi}_h$ collects coefficients on the lagged controls (distinct from the time-series ϕ_h, which stacks the shock and lag-control coefficients into a single vector). The bias of order $O(T^{-1})$ is:

$$E[\hat{\beta}_{\mathrm{LS},N}^h] - \beta^h = -\frac{1}{T-h} \sum_{j=1}^{h} [1 + \varpi_j] \beta^{h-j} + O(T^{-3/2}), \tag{C.5}$$

where $\varpi_j = \frac{1}{N^2} \sum_{i=1}^{N} \sum_{k=1}^{N} \mathrm{tr} \left\{ \Sigma_{c,0,N}^{-1} \Sigma_{c,j,k,i} \right\} \frac{\sigma_{\varepsilon,ik}}{\sigma_{\varepsilon,N}^2}$ collects cross-entity covariance terms; the precise definitions of $\Sigma_{c,0,N}$, $\Sigma_{c,j,k,i}$, $\sigma_{\varepsilon,ik}$, and $\sigma_{\varepsilon,N}^2$ follow Herbst and Johannsen (2024). In large-N panels, the ϖ_j terms become negligible with N^2 scale, and the bias simplifies to:

$$-\frac{1}{T-h} \sum_{j=1}^{h} \beta^{h-j}.$$

The recursive bias-corrected estimator is:

$$\hat{\beta}_{\mathrm{BCC},N}^h = \hat{\beta}_{\mathrm{LS},N}^h + \frac{1}{T-h} \sum_{j=1}^{h} \hat{\beta}_{\mathrm{BCC},N}^{h-j}, \qquad h \geq 1, \tag{C.6}$$

with anchor $\hat{\beta}_{\mathrm{BCC},N}^0 = \hat{\beta}_{\mathrm{LS},N}^0$.

D Bias Reduction by Controlling for Past Information in Local Projections

This appendix explains why adding one lag of the dependent variable as a control in an LPs regression eliminates the future-IRF component of the small-sample bias highlighted in Herbst and Johannsen (2021, 2024). We retain the standard Wold-representation notation θ_ℓ for the moving-average coefficients in this appendix; the LP regression coefficient on the shock at horizon h, denoted β^h in the main text and Appendix C, equals θ_h under the LP regression in equation (D.2).

D.1 Setup

For a covariance-stationary scalar process with Wold representation

$$y_t = \sum_{\ell=0}^{\infty} \theta_\ell \, \varepsilon_{t-\ell}, \qquad \varepsilon_t \overset{\text{i.i.d.}}{\sim} (0, \sigma_\varepsilon^2), \tag{D.1}$$

consider estimating θ_h by the LPs regression

$$y_{t+h} = \alpha^h + \theta_h \, \varepsilon_t + u_t^h, \qquad h \geq 1. \tag{D.2}$$

Without controls, the residual contains all past and future shocks. Herbst and Johannsen (2021) show that this creates bias involving both future (θ_{h+j}) and past (θ_{h-j}) IRF coefficients.

D.2 Adding One Lag of y

We demonstrate that including the lagged outcome truncates the residual to an MA(h) process, thereby eliminating bias arising from serial correlation between the residual and the regressor at horizons beyond h.

Augmenting with y_{t-1}:

$$y_{t+h} = \alpha^h + \theta_h \, \varepsilon_t + \gamma_h \, y_{t-1} + v_t^h. \tag{D.3}$$

Step 1. The raw residual splits into past and future shocks:

$$u_t^h = \underbrace{\sum_{m=1}^{\infty} \theta_{h+m} \, \varepsilon_{t-m}}_{\text{past shocks}} + \underbrace{\sum_{k=1}^{h} \theta_{h-k} \, \varepsilon_{t+k}}_{\text{future shocks}}. \tag{D.4}$$

Step 2. The control $y_{t-1} = \sum_{j=0}^{\infty} \theta_j \, \varepsilon_{t-1-j}$ is a linear combination of the same past shocks.

Step 3. Because shocks are independent, $\text{Cov}(\varepsilon_{t+k}, y_{t-1}) = 0$ for $k \geq 1$, so the projection of u_t^h onto y_{t-1} removes only the past-shock component.

Step 4. The residual after projection:

$$v_t^h = \sum_{k=1}^{h} \theta_{h-k} \, \varepsilon_{t+k}. \tag{D.5}$$

This is an MA(h) error containing only future shocks $\varepsilon_{t+1}, \ldots, \varepsilon_{t+h}$.[10]

Step 5. The bias of the LPs estimator arises from $\mathrm{Cov}(\varepsilon_t, u_t^h)$. Since the future shocks ε_{t+k} ($k \geq 1$) are independent of ε_t by the i.i.d. assumption, only the past-shock component of u_t^h contributes to the bias. By removing this component through the projection onto y_{t-1}, all future-IRF terms θ_{h+j} ($j > 0$) vanish from the bias expression. The remaining bias depends only on past IRF coefficients θ_{h-j}, decaying with T and scaled by the autocovariance of controls.

[10] In an AR(1) process, y_{t-1} is a sufficient statistic for the entire past, so a single lag exactly spans the past-shock component. For higher-order processes, additional lags may be needed to fully remove this component, but a single lag still eliminates the dominant source of bias. This motivates our use of two lags ($p = 2$) in the main specification as an additional safeguard.

E Monte Carlo Simulation Design and Results

We describe the Monte Carlo simulation design used to identify the most suitable LP specification for our setting, with particular attention to the incidental parameter problem. The simulations proceed in three stages: a time-series setup, a panel with an AR(1) process, and a panel with an AR(2) process to verify robustness beyond the baseline autoregressive structure. Specifically, the simulations examine three questions: whether using long differences rather than levels in the outcome variable reduces bias, whether including lagged outcome controls mitigates finite-sample bias as established by Klosin (2024) and Herbst and Johannsen (2024), and whether the analytical bias correction of Herbst and Johannsen (2024) effectively recovers the true impulse response functions.

E.1 Structural Data-Generating Process

We simulate a two-variable structural VAR(1)

$$\boldsymbol{A}_0 \boldsymbol{Y}_t \;=\; \boldsymbol{A}_1 \boldsymbol{Y}_{t-1} + \boldsymbol{\varepsilon}_t, \qquad \boldsymbol{\varepsilon}_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_2), \tag{E.1}$$

with endogenous vector

$\boldsymbol{Y}_t = \big(y_{1t},\, y_{2t}\big)'$, y_{1t} = local economy outcome (e.g. wages, employment), y_{2t} = fishery performance (e.g. revenue).

The contemporaneous structure is upper-triangular,

$$\boldsymbol{A}_0 = \begin{bmatrix} 1 & -0.6 \\ 0 & 1 \end{bmatrix}, \qquad \boldsymbol{A}_0^{-1} = \begin{bmatrix} 1 & 0.6 \\ 0 & 1 \end{bmatrix}, \tag{E.2}$$

thus, a one-unit fishery shock lifts y_{1t} instantaneously by 0.6, while shocks to y_{1t} have no same-period impact on the fishery sector.[11] Lag dynamics start with an AR(1) specification,

$$\boldsymbol{A}_1 = \begin{bmatrix} \rho_1 & \gamma_{21} \\ \gamma_{12} & \phi_{y2} \end{bmatrix} = \begin{bmatrix} \rho_1 & 0.10 \\ 0.10 & 0.50 \end{bmatrix}, \tag{E.3}$$

[11] Using an upper vs. lower triangular $\boldsymbol{A}_0$ only reflects a choice of variable ordering; here the more exogenous series (y_{2t}) is placed last.

featuring weak bidirectional feedback ($\gamma_{12} = \gamma_{21} = 0.10$) and medium persistence of the fishery performance series ($\phi_{y2} = 0.50$). The wage-income persistence ρ_1 drives our three scenarios:

$$\rho_1 \in \{0.30, 0.60, 0.90\} \quad \text{(Low, Medium, High)}.$$

Regarding fishing performance variable, $y_{2,t}$, we assume medium-level persistence with $\phi_{y2} = 0.50$.[12]

Notation. Let $\boldsymbol{Y}_t = (y_{1t}, y_{2t})'$ collect *local economic activity* (y_{1t}, e.g. wages or employment) and *fishery performance* (y_{2t}, e.g. landings or revenue). Shocks $\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I}_2)$ are i.i.d. throughout the time horizon.

E.2 Time-Series SVAR(1)

1. Structural model. The data-generating process (DGP) is a structural VAR(1)

$$\boldsymbol{A}_0 \boldsymbol{Y}_t \;=\; \boldsymbol{A}_1 \boldsymbol{Y}_{t-1} + \boldsymbol{\varepsilon}_t, \qquad \boldsymbol{A}_0 \;=\; \begin{bmatrix} 1 & -0.6 \\ 0 & 1 \end{bmatrix}, \quad \boldsymbol{A}_0^{-1} = \begin{bmatrix} 1 & 0.6 \\ 0 & 1 \end{bmatrix}. \tag{E.4}$$

The triangular structure allows a fishery-specific shock (ε_{2t}) to contemporaneously increase y_{1t} by 0.6, while ruling out any immediate feedback from y_{1t} to the fishery sector. Notably, the inverse of $\boldsymbol{A}_0$ is key to identifying the exogenous impulse-response structure from fisheries to the local economy.

2. Lag dynamics. The single-lag matrix that has information about auto-regressive persistence of each variable and cross-impact is

$$\boldsymbol{A}_1 = \begin{bmatrix} \rho_1 & \gamma_{21} \\ \gamma_{12} & \phi_{y2} \end{bmatrix} = \begin{bmatrix} \rho_1 & 0.10 \\ 0.10 & 0.50 \end{bmatrix},$$

where only the persistence of y_{1t} varies in order to investigate how persistence of the autoregressive process governs bias in IRF estimation:

$$\rho_1 \in \{0.30,\ 0.60,\ 0.90\} \quad \text{(Low, Medium, High)}.$$

3. Monte Carlo settings. $R = 500$ replications, $T = 16$ (replicating our limited finite time sample), $H = 6$ for IRF estimation horizon, initial state $(y_{10}, y_{20}) = (10, 5)$.

[12] We tested the sensitivity of different persistence cases of the fishing variable. We found no visible change in results.

4. Population impulse response. Let $\boldsymbol{B}_1 = \boldsymbol{A}_0^{-1}\boldsymbol{A}_1$ and define the moving-average coefficient matrix $\boldsymbol{\Phi}_h = \boldsymbol{B}_1^h$. For a unit structural fishery shock the true IRF is

$$\mathrm{IRF}_h = \left(\boldsymbol{\Phi}_h \, \boldsymbol{A}_0^{-1}\right)_{1,2}, \qquad h = 0, \ldots, H.$$

5. Local-projection estimators under multiple specifications. Following Klosin (2024) and Herbst and Johannsen (2024) on the role of lagged outcome controls in reducing finite-sample IRF bias, Olea and Plagborg-Møller (2021) on lag augmentation for robust inference, and Piger and Stockwell (2025, 2023) on the bias-reduction properties of long-difference specifications, we test four specifications. For each horizon h, four LPs regressions are run on the same simulated series:

1. **Level + Lag controls (i.e. past information)**
   $y_{1,t+h} = \alpha^h + \beta^h \, y_{2t} + [y_{1,t-1},\, y_{1,t-2},\, y_{2,t-1},\, y_{2,t-2}]\, \Gamma^h + u_t^h$

2. **Level (No controls)**
   $y_{1,t+h} = \alpha^h + \beta^h \, y_{2t} + u_t^h$

3. **Long-difference + Lag controls**
   $\Delta^h y_{1,t} = (y_{1,t+h} - y_{1,t-1}) = \alpha^h + \beta^h \, y_{2t} + [y_{1,t-1},\, y_{1,t-2},\, y_{2,t-1},\, y_{2,t-2}]\, \Pi^h + u_t^h$

4. **Long-difference (No controls)**
   $\Delta^h y_{1,t} = (y_{1,t+h} - y_{1,t-1}) = \alpha^h + \beta^h \, y_{2t} + u_t^h.$

Storing $\widehat{\beta}^h$ over R replications yields the MC mean, point-wise 95% envelope, and bias ($\mathrm{Bias}^h = \mathbb{E}[\widehat{\beta}^h] - \mathrm{IRF}_h$).

E.3 Panel SVAR(1)

1. Structural model. For $i = 1, \ldots, N$ individuals

$$\boldsymbol{A}_0 \, \boldsymbol{Y}_{i,t} \;=\; \boldsymbol{\alpha}_i + \boldsymbol{A}_1 \, \boldsymbol{Y}_{i,t-1} + \boldsymbol{\varepsilon}_{i,t}, \qquad \boldsymbol{\varepsilon}_{i,t} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_2), \quad \boldsymbol{\alpha}_i \sim \mathcal{N}\big((3,3)', \, 2^2 \boldsymbol{I}_2\big). \tag{E.5}$$

The matrices $\boldsymbol{A}_0$ and $\boldsymbol{A}_1$ are identical to the time-series case; only ρ_1 varies by scenario. In the panel DGP, we incorporate community-fixed effects for each community i and up to $N = 150$ fishing communities, in order to examine the bias resulting from IPP in the dynamic panel setup with pre-determined persistence parameter of the local economy variable.

2. Monte Carlo settings. [13]

$$R = 300, \quad N = 150, \quad T = 16, \quad H = 6.$$

3. True IRF. Because model parameters are common across i, the population IRF again equals $\left(\mathbf{\Phi}_h \, \boldsymbol{A}_0^{-1}\right)_{1,2}$ with $\mathbf{\Phi}_h = \boldsymbol{B}_1^h$.

4. Estimation. For each h and replication we fit four panel LPs with individual fixed effects (demeaning via `felm`):

$$\text{specifications (1)–(4) as in Section E.2,} \quad \text{but on panel data and with } \alpha_i^h.$$

MC statistics are computed exactly as in the time-series study.

E.4 Panel SVAR(2)

Structural form (individual i, time t).

$$\boldsymbol{A}_0 \boldsymbol{Y}_{i,t} \;=\; \boldsymbol{\alpha}_i \;+\; \boldsymbol{A}_1 \boldsymbol{Y}_{i,t-1} \;+\; \boldsymbol{A}_2 \boldsymbol{Y}_{i,t-2} \;+\; \boldsymbol{\varepsilon}_{i,t}, \qquad \boldsymbol{\varepsilon}_{i,t} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \boldsymbol{I}_2), \tag{E.6}$$

where all previous DGP specifications are identically applied except for the lag structure with AR(2) process. The fixed effects $\boldsymbol{\alpha}_i$ follow the same distribution as in Section E.3.

Contemporaneous matrix (unchanged).

$$\boldsymbol{A}_0 = \begin{bmatrix} 1 & -0.6 \\ 0 & 1 \end{bmatrix}, \quad \boldsymbol{A}_0^{-1} = \begin{bmatrix} 1 & 0.6 \\ 0 & 1 \end{bmatrix}.$$

Lag matrices.

$$\boldsymbol{A}_1 = \begin{bmatrix} \rho_1 & \gamma_{21} \\ \gamma_{12} & \phi_{y2} \end{bmatrix}, \qquad \boldsymbol{A}_2 = \begin{bmatrix} \rho_2 & \delta_{21} \\ \delta_{12} & 0 \end{bmatrix},$$

with fixed cross-lag parameters $\gamma_{21} = \gamma_{12} = +0.10$, $\delta_{21} = \delta_{12} = -0.05$, $\phi_{y2} = 0.50$.

[13] We limit the number of replications to $R = 300$ because the simulation with individual fixed effects requires considerable computational time.

In the AR(2) process, the sum of autoregressive coefficients determines the general data persistence and stability (i.e., $|\rho_1 + \rho_2|$). Three persistence scenarios are considered:

$$(\rho_1, \rho_2) : \text{Low } = (0.40, -0.10), \qquad \text{Medium } = (0.45, 0.15), \qquad \text{High } = (0.60, 0.30).$$

True impulse response. Write $\boldsymbol{B}_1 = \boldsymbol{A}_0^{-1}\boldsymbol{A}_1$ and $\boldsymbol{B}_2 = \boldsymbol{A}_0^{-1}\boldsymbol{A}_2$. For a VAR(2) the moving-average coefficient matrices satisfy $\boldsymbol{\Phi}_0 = \boldsymbol{I}_2$, $\boldsymbol{\Phi}_1 = \boldsymbol{B}_1$, $\boldsymbol{\Phi}_h = \boldsymbol{B}_1\boldsymbol{\Phi}_{h-1} + \boldsymbol{B}_2\boldsymbol{\Phi}_{h-2}$ $(h \geq 2)$. The population IRF to a one-unit fishery shock is

$$\text{IRF}_h \;=\; \left(\boldsymbol{\Phi}_h\, \boldsymbol{A}_0^{-1}\right)_{1,2}, \qquad h = 0, \ldots, H.$$

The rest of the procedure is equivalent to the time-series and panel VAR(1) parts.

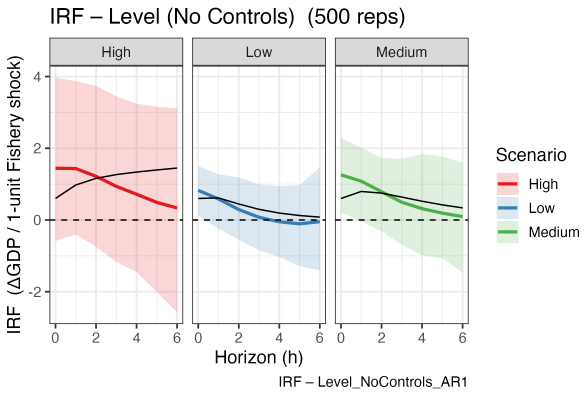
Output. For every scenario-specification pair we report:

$$\left\{ \overline{\beta}^h,\ 95\%\text{ envelope},\ \text{Bias}^h = \overline{\beta}^h - \text{IRF}_h \right\}, \quad h = 0, \ldots, H,$$

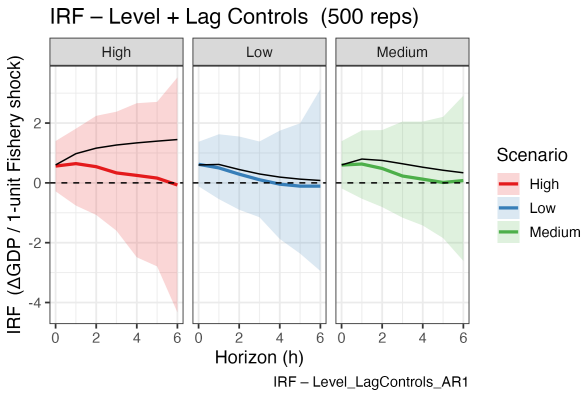
visualized as shaded IRF bands and corresponding bias plots.

E.5 Simulation Results

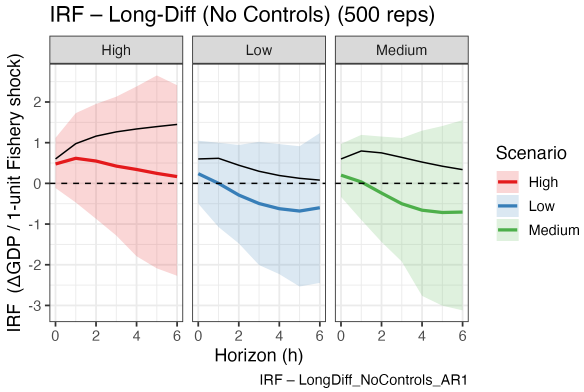
Time-Series



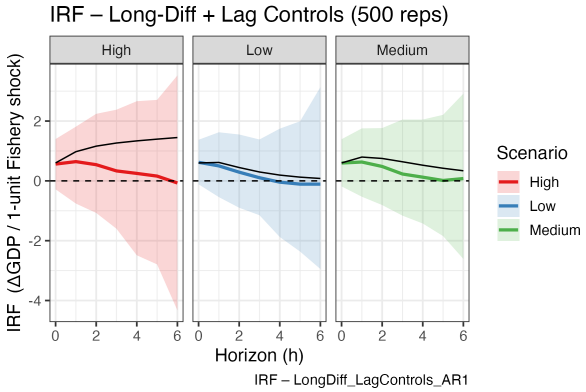
(a) IRF: Levels, No Controls



(b) IRF: Levels, Baseline



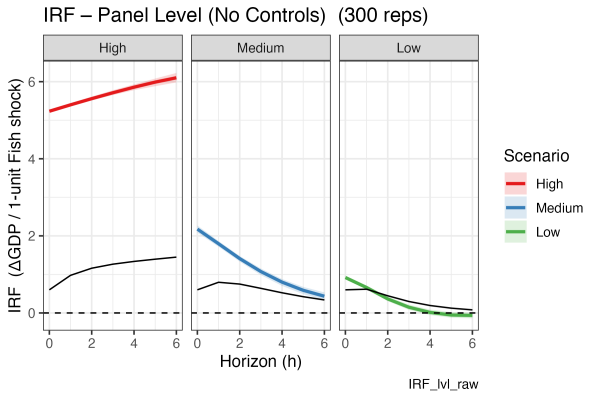
(c) IRF: Long Differences, No Controls



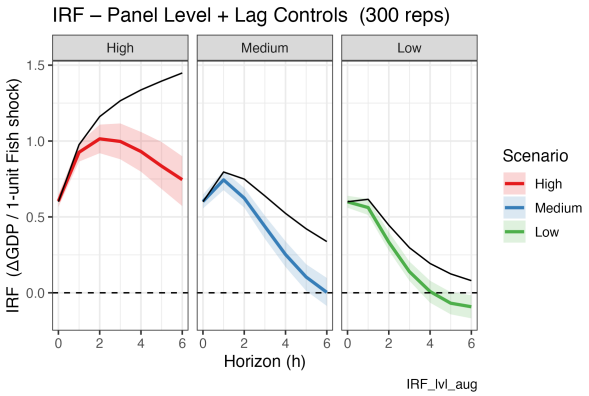
(d) IRF: Long Differences, Baseline

Figure E.1: Impulse response functions (IRFs) estimated under different specifications: levels vs. long-difference outcomes, with and without baseline controls ($p = q = 2$ lags of outcome and treatment).

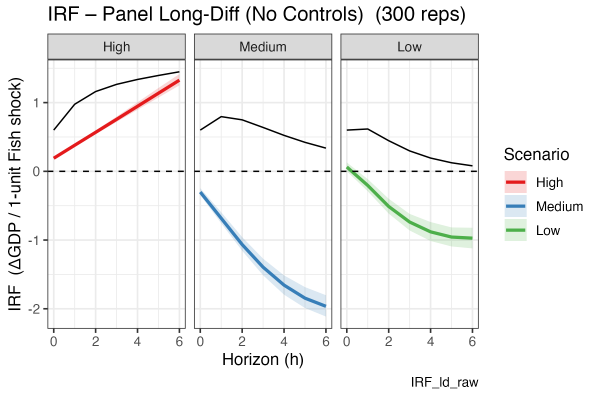
Panel AR(1)



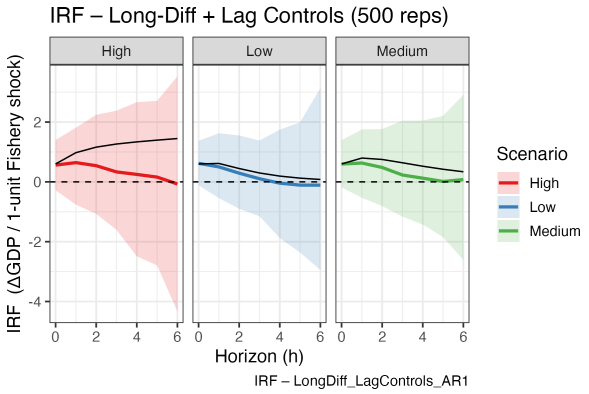
(a) Panel IRF: Levels, No Controls



(b) Panel IRF: Levels, Baseline



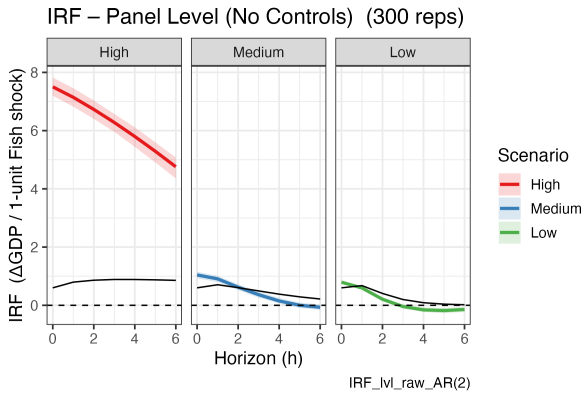
(c) Panel IRF: Long Differences, No Controls



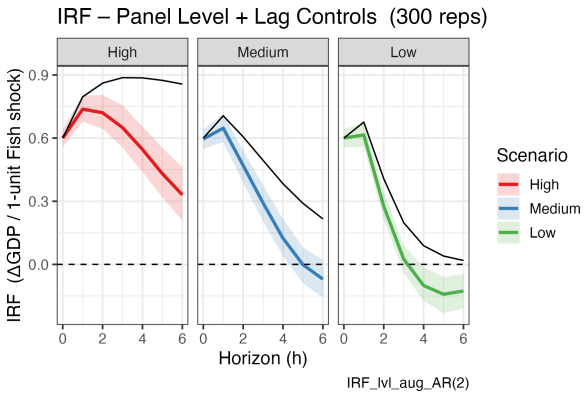
(d) Panel IRF: Long Differences, Baseline

Figure E.2: Panel impulse response functions (IRFs) across different specifications: levels vs. long differences, with and without baseline controls ($p = q = 2$).

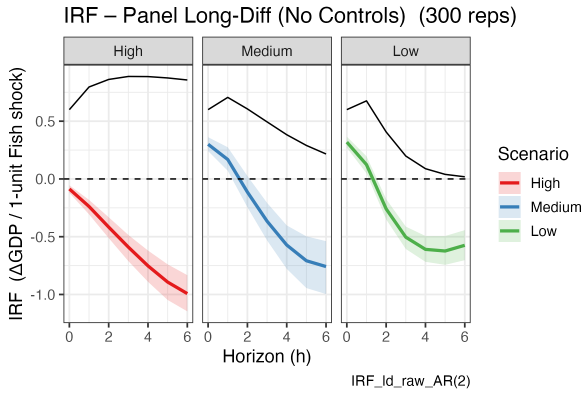
Panel AR(2)



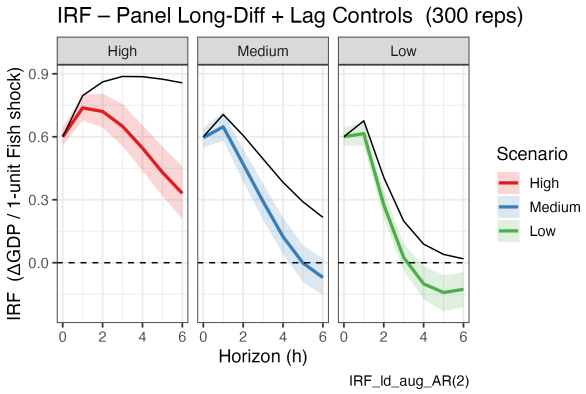
(a) Panel IRF (AR(2)): Levels, No Controls



(b) Panel IRF (AR(2)): Levels, Baseline



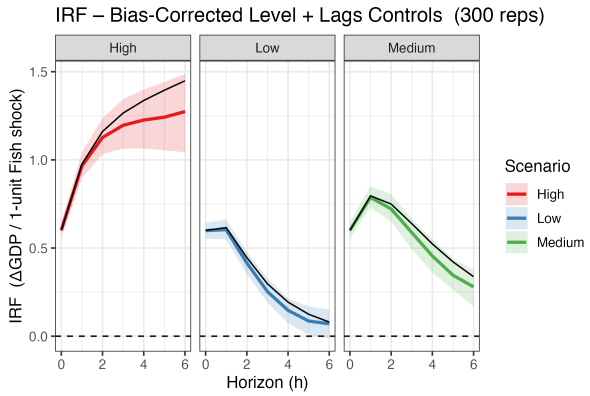
(c) Panel IRF (AR(2)): Long Differences, No Controls



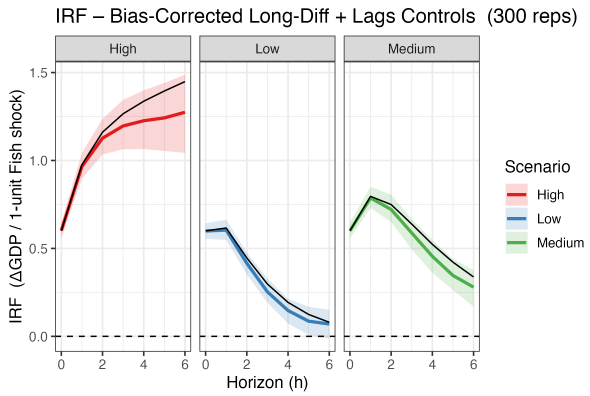
(d) Panel IRF (AR(2)): Long Differences, Baseline

Figure E.3: Panel impulse response functions (IRFs) under an AR(2) specification: levels vs. long differences, with and without baseline controls ($p = q = 2$).

Bias-Corrected Results



(a) Bias-corrected IRF: Levels, Baseline



(b) Bias-corrected IRF: Long Differences, Baseline

Figure E.4: Bias-corrected impulse response functions (IRFs) with baseline controls ($p = q = 2$), following Herbst and Johannsen (2021, 2024).

F System GMM Estimation for Joint Inference

To conduct joint inference on the impulse response path, we estimate the IRFs as a GMM problem following Jordà and Taylor (2025). The implementation requires estimating the cross-horizon covariance matrix of estimated IRFs, $\hat{\Omega}_\beta$, which in turn requires all horizon-specific estimates to be computed on a common sample—observations for which all outcomes $\Delta^h y_{i,t+h}$ for $h = 0, \ldots, 5$ are simultaneously non-missing. This common sample restriction reduces the effective sample size by approximately 40% at shorter horizons relative to the horizon-specific regressions, which inevitably may affect point estimate values, with larger confidence intervals.

No existing statistical package simultaneously accommodates panel fixed effects, instrumental variables, cluster-robust standard errors, and multi-equation GMM covariance estimation. We therefore employ a Frisch-Waugh-Lovell (FWL) residualization procedure, partialling out lag controls and fixed effects before computing the IV estimator from the residualized variables. The FWL theorem guarantees numerical equivalence with direct estimation; we verify this by confirming that the GMM and individual LP-IV estimates nearly coincide at $h = 5$, where both use effectively the same sample.[14]

Figures F.1 and F.2 compare the individual LP-IV and system GMM point estimates. Differences arise solely from sample composition; the two converge at $h = 5$ where samples coincide and preserve the same qualitative trajectory throughout, confirming the validity of the GMM approach.

[14] Small residual differences at $h = 5$ may arise from numerical precision in the optimization procedure and minor differences in how the statistical packages handle sample selection and weighting.

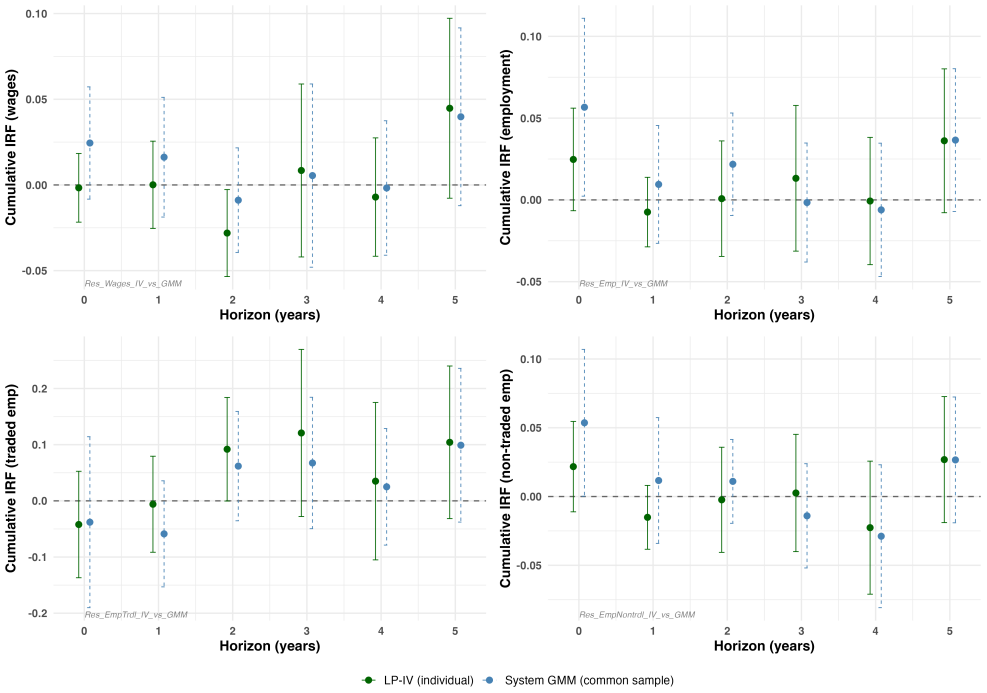


Figure F.1: Comparison of individual LP-IV and system GMM point estimates for resident-based outcomes. Green dots denote LP-IV estimates using horizon-specific samples; blue dots denote system GMM estimates using the common sample. Vertical bars represent 90% confidence intervals.

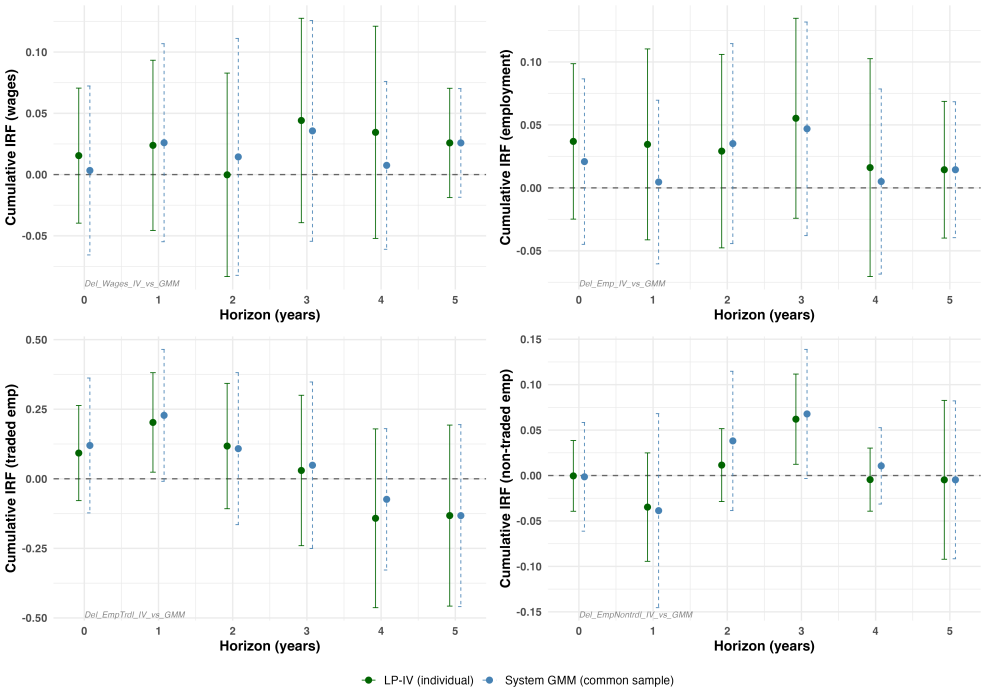


Figure F.2: Comparison of individual LP-IV and system GMM point estimates for delivery-based outcomes. Green dots denote LP-IV estimates using horizon-specific samples; blue dots denote system GMM estimates using the common sample. Vertical bars represent 90% confidence intervals.

G Summary Statistics

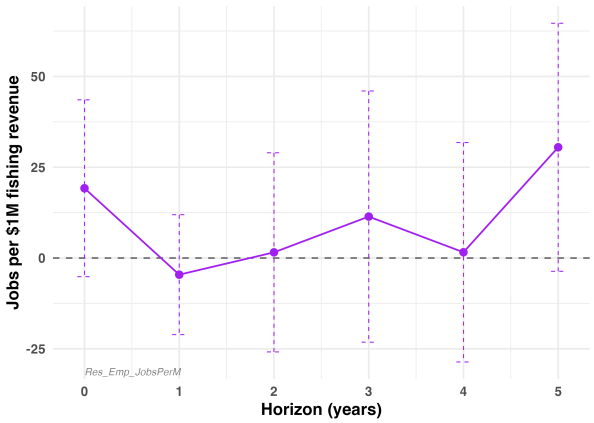
	Resident earnings			Local landings		
Variable	N	Mean	SD	N	Mean	SD
Fishing revenue ($)	1,687	2,250,133	7,353,771	577	10,450,062	15,047,562
Total wages ($)	1,687	33,934,174	243,635,258	577	84,553,268	410,988,485
Total employment	1,687	1,754	11,631	577	4,232	19,585
Traded employment	1,687	126	645	577	278	1,073
Non-traded employment	1,687	1,628	11,001	577	3,955	18,529
Communities		148			48	

All monetary values are in real (inflation-adjusted) dollars. Employment is measured as the number of individual workers (headcount).

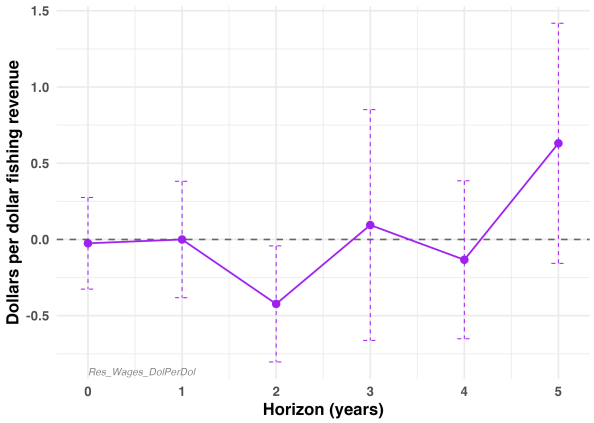
Table G.1: Summary statistics for the estimation sample at $h = 0$ (baseline specification).

H Scaled Impulse Response Functions

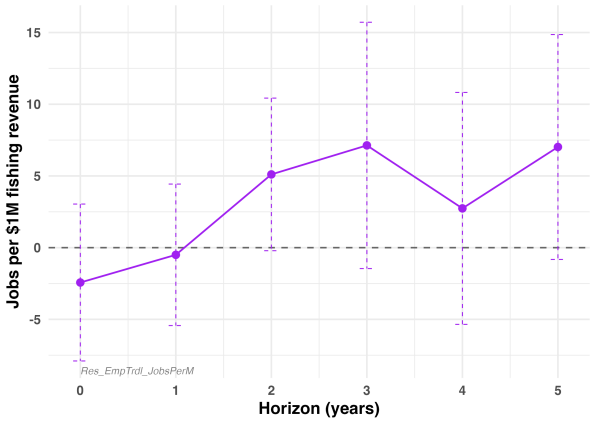
This appendix presents the estimated impulse response functions converted from elasticities to level effects using the transformation $\Delta Y/\$ = \hat{\beta}^h_{\text{BCC}} \cdot \bar{y}/\bar{x}$. Employment effects are expressed as jobs per million dollars of fishing revenue; wage effects are expressed as dollars per dollar of fishing revenue. Dashed error bars indicate 90% confidence intervals.



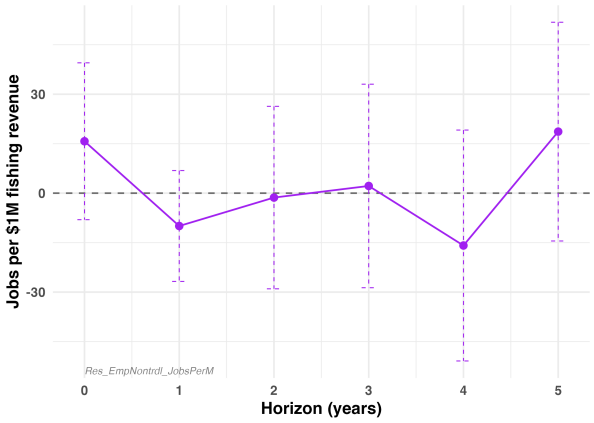
(a) Total employment



(b) Wages

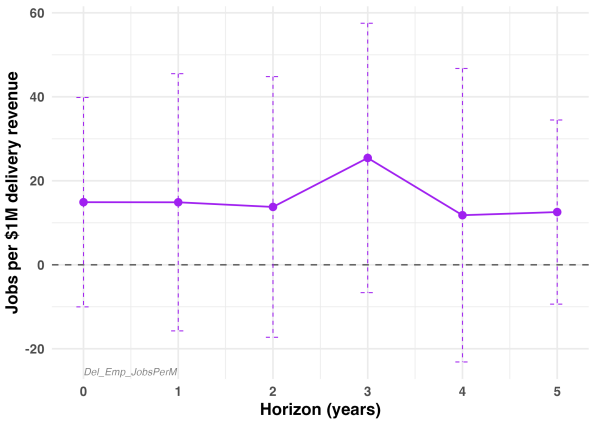


(c) Traded employment

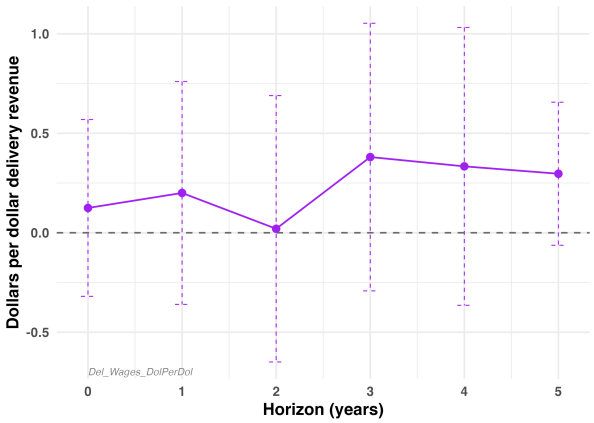


(d) Non-traded employment

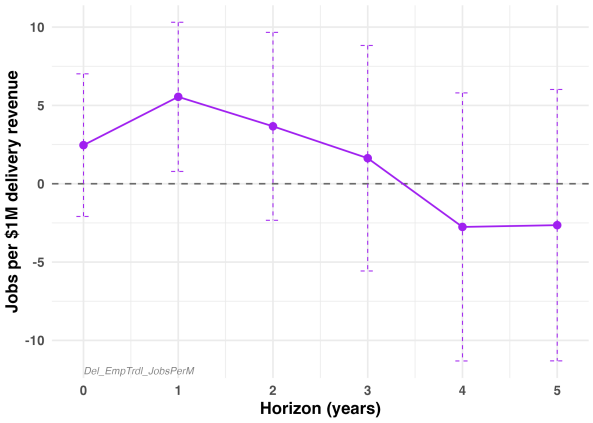
Figure H.1: Scaled cumulative impulse response functions, resident specification (LP-IV-BCC).



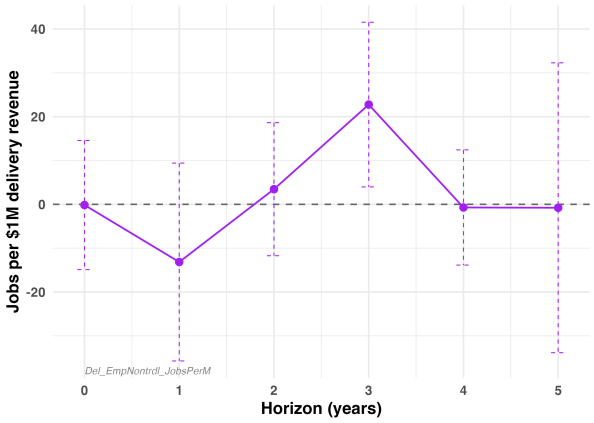
(a) Total employment



(b) Wages



(c) Traded employment



(d) Non-traded employment

Figure H.2: Scaled cumulative impulse response functions, delivery specification (LP-IV-BCC).

Chapter 5

Conclusion

This dissertation has examined how coastal communities and the marine systems they depend on interact under environmental and economic uncertainty, drawing on three essays that span structural economics, dynamic bioeconomic modeling, and causal panel inference. Taken together, the essays show that coastal community resilience under environmental change depends on the joint configuration of three complementary channels—the *ex-ante* structure of local economic portfolios (Chapter 2), the *active management* of degraded ecosystems (Chapter 3), and the *ex-post* propagation of shocks through local economies (Chapter 4). No single channel is sufficient on its own; understanding coastal resilience requires considering how these dimensions interact. As climate-driven change continues to reshape coastal and marine systems, sustained engagement with the coupled human-environmental systems perspective developed here will remain essential to the design of policies that support both ecological function and the welfare of communities that depend on marine resources.

Bibliography

J. K. Abbott, Y. Sakai, and D. S. Holland. Species, space and time: A quarter century of fishers' diversification strategies on the US West Coast. *Fish and Fisheries*, 24:93–110, 1 2023. ISSN 14672979. doi: 10.1111/faf.12712.

D. Aigner, C. A. Lovell, and P. Schmidt. Formulation and estimation of stochastic frontier production function models. *Journal of Econometrics*, 6:21–37, 7 1977. ISSN 0304-4076. doi: 10.1016/0304-4076(77)90052-5.

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.

H. Allcott and D. Keniston. Dutch disease or agglomeration? the local economic effects of natural resource booms in modern america. *The Review of Economic Studies*, 85(2):695–731, 4 2018. ISSN 0034-6527. doi: 10.1093/restud/rdx042.

J. Alvarez and M. Arellano. The time series and cross-section asymptotics of dynamic panel data estimators. *Econometrica*, 71:1121–1159, 2003. doi: 10.1111/1468-0262.00441.

C. Amsler, A. Prokhorov, and P. Schmidt. Endogeneity in stochastic frontier models. *Journal of Econometrics*, 190:280–288, 2 2016. ISSN 18726895. doi: 10.1016/j.jeconom.2015.06.013.

J. L. Anderson, F. Asche, and T. Garlock. Globalization and commoditization: The transformation of the seafood market. *Journal of Commodity Markets*, 12:2–8, 12 2018. ISSN 2405-8513. doi: 10.1016/J.JCOMM.2017.12.004.

S. C. Anderson, E. J. Ward, A. O. Shelton, M. D. Adkison, A. H. Beaudreau, R. E. Brenner, A. C. Haynie, J. C. Shriver, J. T. Watson, and B. C. Williams. Benefits and risks of diversification for individual fishers. *Proceedings of the National Academy of Sciences of the United States of America*, 114:10797–10802, 10 2017. ISSN 10916490. doi: 10.1073/pnas.1702506114.

R. Arezki, V. A. Ramey, and L. Sheng. News shocks in open economies: Evidence from giant oil discoveries. *The Quarterly Journal of Economics*, 132(1):103–155, 2017.

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.

ASMI. The economic value of alaska's seafood industry. Technical report, McKinley Research Group, April 2024. URL `https://www.alaskaseafood.org/wp-content/uploads/MRG_ASMI-Economic-Impacts-Report_2023_WEB-PAGES.pdf`. Accessed: 2024-11-30.

S. E. Atkinson, R. Färe, and D. Primont. Stochastic estimation of firm inefficiency using distance functions. *Southern Economic Journal*, 69:596–611, 1 2003. ISSN 0038-4038. doi: 10.1002/j.2325-8012.2003.tb00515.x.

A. J. Auerbach and Y. Gorodnichenko. Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, 4:1–27, 5 2012. doi: 10.1257/pol.4.2.1.

M. Baron, E. Verner, and W. Xiong. Banking crises without panics. *Quarterly Journal of Economics*, 136: 51–113, 2 2021. doi: 10.1093/qje/qjaa034.

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.

F. Belotti, S. Daidone, G. Ilardi, and V. Atella. Stochastic frontier analysis using stata. *The Stata Journal*, 13:719–758, 2013.

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.

A. M. Birkenbach, A. L. Cojocaru, and R. B. M. Pincinato. Coastal communities at a crossroads. *Marine Resource Economics*, 38(4):353–364, 9 2023.

R. Blasiak, J. Spijkers, K. Tokunaga, J. Pittman, N. Yagi, and H. Österblom. Climate change and marine

fisheries: Least developed countries top global index of vulnerability. *PLoS ONE*, 12, 6 2017. ISSN 19326203. doi: 10.1371/journal.pone.0179632.

K. Borusyak, P. Hull, and X. Jaravel. A practical guide to shift-share instruments. *NBER Working Paper*, 12 2024. doi: 10.3386/W33236.

A. L. Bris, K. E. Mills, R. A. Wahle, Y. Chen, M. A. Alexander, A. J. Allyn, J. G. Schuetz, J. D. Scott, and A. J. Pershing. Climate vulnerability and resilience in the most valuable North American fishery. *Proceedings of the National Academy of Sciences of the United States of America*, 115:1831–1836, 2 2018. ISSN 10916490. doi: 10.1073/pnas.1711122115.

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.

B. Brümmer, T. Glauben, and G. Thijssen. Decomposition of productivity growth using distance functions: The case of dairy farms in three European countries. *American Journal of Agricultural Economics*, 84: 628–644, 8 2002. ISSN 14678276. doi: 10.1111/1467-8276.00324.

N. P. Burnett, A. M. Ricart, T. Winquist, A. M. Saley, M. S. Edwards, B. Hughes, J. Hodin, M. L. Baskett, and B. Gaylord. Bimodal spore release heights in the water column enhance local retention and population connectivity of bull kelp, nereocystis luetkeana. *Ecology and Evolution*, 14, 8 2024. ISSN 20457758. doi: 10.1002/ece3.70177.

California Department of Fish and Wildlife. Red sea urchin, *Mesocentrotus franciscanus*, enhanced status report, 2020. Available at: `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. Technical report, California Ocean Protection Council, Sacramento, 2 2021. Contributors: Michael Esgro and James Ray. Available at: `https://www.opc.ca.gov/webmaster/ftp/pdf/agenda_items/20210216/Item7_KelpActionPlan_ExhibitA_FINAL.pdf`.

P. E. Carnell, K. Whiteoak, M. Young, K. Critchell, S. Swearer, P. I. Macreadie, J. McIntyre, and E. A.

Treml. 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. Carothers. Tragedy of commodification: Transitions in Alutiiq fishing communities in the Gulf of Alaska. *Maritime Studies (MAST)*, 9:95–120, 2010.

G. Chamberlain. Multivariate regression models for panel data. *Journal of Econometrics*, 18:5–46, 1 1982. ISSN 0304-4076. doi: 10.1016/0304-4076(82)90094-X.

S. Chandra. A test of the regional growth-instability frontier using state data. *Land Economics*, 78:442–462, 2002.

S. Chandra. Regional economy size and the growth instability frontier: Evidence from Europe. *Journal of Regional Science*, 43:95–122, 2003. ISSN 00224146. doi: 10.1111/1467-9787.00291.

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.

P. M. Clay and J. Olson. Defining "fishing communities": Vulnerability and the Magnuson-Stevens fishery conservation and management act. *Human Ecology Review*, 15:143–160, 2008.

T. J. Cline, D. E. Schindler, and R. Hilborn. Fisheries portfolio diversification and turnover buffer Alaskan fishing communities from abrupt resource and market changes. *Nature Communications*, 8, 1 2017. ISSN 20411723. doi: 10.1038/ncomms14042.

T. Coelli. On the econometric estimation of the distance function representation of a production technology. *LIDAM Discussion Papers CORE*, 9 2000. URL `https://ideas.repec.org/p/cor/louvco/2000042.html`.

T. Coelli and S. Perelman. Efficiency measurement, multi-output technologies and distance functions: With application to European railways. Working Paper 96/05, Centre de Recherche en Economie Publique Et

Economie de la Population, Université de Liège, Liège, Belgium, 1996. URL http://orbi.ulg.be/bitstream/2268/35665/1/105.%20CREPP%209605%20Coelli-Perelman.pdf.

L. L. Colburn, M. Jepson, C. Weng, T. Seara, J. Weiss, and J. A. Hare. Indicators of climate change and social vulnerability in fishing dependent communities along the Eastern and Gulf Coasts of the United States. *Marine Policy*, 74:323–333, 12 2016. ISSN 0308597X. doi: 10.1016/j.marpol.2016.04.030.

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.

R. A. Cuesta and L. Orea. Mergers and technical efficiency in Spanish savings banks: A stochastic distance function approach. *Journal of Banking and Finance*, 26:2231–2247, 1 2002. ISSN 0378-4266. doi: 10.1016/S0378-4266(01)00184-4.

R. A. Cuesta and J. Zofío. Hyperbolic efficiency and parametric distance functions: With application to Spanish savings banks. *Journal of Productivity Analysis*, 24:31–48, 9 2005.

R. A. Cuesta, C. A. Lovell, and J. L. Zofío. Environmental efficiency measurement with translog distance functions: A parametric approach. *Ecological Economics*, 68:2232–2242, 6 2009. ISSN 09218009. doi: 10.1016/j.ecolecon.2009.02.001.

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.

B. Dalheimer, I. Parikoglou, F. Brambach, M. Yanita, H. Kreft, and B. Brümmer. On the palm oil-biodiversity trade-off: Environmental performance of smallholder producers. *Journal of Environmental Economics and Management*, 125, 5 2024. ISSN 10960449. doi: 10.1016/j.jeem.2024.102975.

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 2001. 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. Available at: `https://sites.google.com/site/mdpsolve/users-guide`.

FAO. *The State of World Fisheries and Aquaculture 2024 — Blue Transformation in action*. Food and Agriculture Organization of the United Nations, Rome, Italy, 2024. URL `https://www.fao.org/publications/sofia/2024`.

M. C. Fisher, S. K. Moore, S. L. Jardine, J. R. Watson, and J. F. Samhouri. Climate shock effects and mediation in fisheries. *Proceedings of the National Academy of Sciences*, 118(2):e2014379117, 1 2021.

K. Fishman. Red urchin fishery still struggling on North Coast after kelp collapse; 'everybody's scared', Aug. 2022. The Mendocino Voice. Available at: `https://mendovoice.com/2022/08/red-urchin-fishery-still-struggling-in-north-coast-after-kelp-collapse-everybodys-scared/` (accessed 2023-08-06).

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. M. Free, J. T. Thorson, M. L. Pinsky, K. L. Oken, J. Wiedenmann, and O. P. Jensen. Impacts of historical warming on marine fisheries production. *Science*, 363(6430):979–983, 3 2019.

S. Gokhale, J. Blomquist, M. Lindegren, A. Richter, and S. Waldo. The role of non-fishing and partner incomes in managing fishers' economic risk. *Marine Resource Economics*, 39:285–305, 10 2024. ISSN 23345985. doi: 10.1086/731762.

W. Greene. Fixed and random effects in stochastic frontier models. *Journal of Productivity Analysis*, 23: 7–32, 2005a.

W. Greene. Reconsidering heterogeneity in panel data estimators of the stochastic frontier model. *Journal of Econometrics*, 126:269–303, 6 2005b. ISSN 03044076. doi: 10.1016/j.jeconom.2004.05.003.

W. E. Griffiths and G. Hajargasht. Some models for stochastic frontiers with endogeneity. *Journal of Econometrics*, 190:341–348, 2 2016. ISSN 18726895. doi: 10.1016/j.jeconom.2015.06.012.

M. Guettabi and A. James. Who benefits from an oil boom? evidence from a unique Alaskan data set. *Resource and Energy Economics*, 62:101200, 11 2020. ISSN 0928-7655. doi: 10.1016/J.RESENEECO.2020.101200.

K. A. Hafner. Growth-instability frontier and industrial diversification: Evidence from European gross value added. *Papers in Regional Science*, 98:799–824, 4 2019. ISSN 14355957. doi: 10.1111/pirs.12401.

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.

E. P. Herbst and B. K. Johannsen. Bias in local projections. *Finance and Economics Discussion Series*, 2020:1–62, 1 2021. doi: 10.17016/FEDS.2020.010R1.

E. P. Herbst and B. K. Johannsen. Bias in local projections. *Journal of Econometrics*, 240, 3 2024. doi: 10.1016/j.jeconom.2024.105655.

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.

A. Himes-Cornell and K. Hoelting. Resilience strategies in the face of short- and long-term change: Out-migration and fisheries regulation in Alaskan fishing communities. *Ecology and Society*, 20, 6 2015. ISSN 17083087. doi: 10.5751/ES-07074-200209.

A. Himes-Cornell and S. Kasperski. Assessing climate change vulnerability in Alaska's fishing communities. *Fisheries Research*, 162:1–11, 2 2015. ISSN 01657836. doi: 10.1016/j.fishres.2014.09.010.

A. J. Hobday, L. V. Alexander, S. E. Perkins, D. A. Smale, S. C. Straub, E. C. Oliver, J. A. Benthuysen, 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.

D. S. Holland and S. Kasperski. The impact of access restrictions on fishery income diversification of US West Coast fishermen. *Coastal Management*, 44:452–463, 9 2016. ISSN 15210421. doi: 10.1080/08920753.2016.1208883.

D. S. Holland, C. Speir, J. Agar, S. Crosson, G. Depiper, S. Kasperski, A. W. Kitts, and L. Perruso. Impact of catch shares on diversification of fishers' income and risk. *Proceedings of the National Academy of Sciences of the United States of America*, 114:9302–9307, 8 2017. ISSN 10916490. doi: 10.1073/pnas.1702382114.

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.

G. D. Jacobsen, D. P. Parker, and J. B. Winikoff. Are resource booms a blessing or a curse?: Evidence from people (not places). *Journal of Human Resources*, 58(2):393–420, 2023.

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.

O. Jordà. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, 3 2005. doi: 10.1257/0002828053828518.

O. Jordà. Local projections for applied economics. *Annual Review of Economics*, 15:607–631, 9 2023. doi: 10.1146/ANNUREV-ECONOMICS-082222-065846.

Ò. Jordà and A. M. Taylor. The time for austerity: Estimating the average treatment effect of fiscal policy. *Economic Journal*, 126:219–255, 2 2016. ISSN 14680297. doi: 10.1111/ecoj.12332.

O. Jordà and A. M. Taylor. Local projections. *Journal of Economic Literature*, 63:59–110, 3 2025. doi: 10.1257/jel.20241521.

O. Jordà, M. Schularick, and A. M. Taylor. Betting the house. *Journal of International Economics*, 96: S2–S18, 7 2015. doi: 10.1016/J.JINTECO.2014.12.011.

O. Jordà, M. Schularick, and A. M. Taylor. The effects of quasi-random monetary experiments. *Journal of Monetary Economics*, 112:22–40, 6 2020. doi: 10.1016/J.JMONECO.2019.01.021.

L. Jost. Entropy and diversity. *Oikos*, 113:363–375, 5 2006. ISSN 00301299. doi: 10.1111/j.2006.0030-1299.14714.x.

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. Kasperski and D. S. Holland. Income diversification and risk for fishermen. *Proceedings of the National Academy of Sciences of the United States of America*, 110:2076–2081, 2 2013. ISSN 00278424. doi: 10.1073/pnas.1212278110.

S. Klosin. Dynamic biases of static panel data estimators. 10 2024. URL https://arxiv.org/abs/2410.16112v1.

J. Kluge. Sectoral diversification as insurance against economic instability. *Journal of Regional Science*, 58: 204–223, 1 2018. ISSN 14679787. doi: 10.1111/jors.12349.

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.

S. C. Kumbhakar, H.-J. Wang, and A. P. Horncastle. *A Practitioner's Guide to Stochastic Frontier Analysis Using Stata*. Cambridge University Press, 1 2015. ISBN 9781139342070. doi: 10.1017/CBO9781139342070.

P. S. Lande. Regional industrial structure and economic growth and instability. *Journal of Regional Science*, 34:343–360, 8 1994. ISSN 1467-9787. doi: 10.1111/J.1467-9787.1994.TB00871.X.

A. Lavoie and A. Himes-Cornell. Social networks of Alaska fishing communities. *Coastal Management*, 47:1–22, 1 2019. ISSN 15210421. doi: 10.1080/08920753.2019.1525259.

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.

H. Markowitz. Portfolio selection. *The Journal of Finance*, 7:77–91, 1952. ISSN 15406261. doi: 10.1111/J.1540-6261.1952.TB01525.X.

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. Mian, A. Sufi, and E. Verner. Household debt and business cycles worldwide. *Quarterly Journal of Economics*, 132:1755–1817, 11 2017. doi: 10.1093/qje/qjx017.

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.

E. Moretti. Local multipliers. *American Economic Review*, 100:373–377, 5 2010. ISSN 00028282. doi: 10.1257/aer.100.2.373.

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.

C. J. Morrison-Paul, W. E. Johnston, and G. A. Frengley. Efficiency in New Zealand sheep and beef farming: The impacts of regulatory reform. *Review of Economics and Statistics*, 82:325–337, 2000. ISSN 00346535. doi: 10.1162/003465300558713.

Y. Mundlak. On the pooling of time series and cross section data. *Econometrica*, 46:69, 1 1978. ISSN 00129682. doi: 10.2307/1913646.

S. Nickell. Biases in dynamic models with fixed effects. *Econometrica*, 49:1417–1426, 11 1981. doi: 10.2307/1911408.

NOAA. Fisheries of the united states, 2020. Technical report, NOAA Current Fishery Statistics, May 2022. URL `https://media.fisheries.noaa.gov/2022-05/Fisheries-of-the-United-States-2020-Report-FINAL.pdf`. Accessed: 2022-05-01.

NOAA Office for Coastal Management. 2024 marine economy report: U.s. territories. Marine economy report, National Oceanic and Atmospheric Administration, 2024. URL `https://coast.noaa.gov/digitalcoast/data/enow.html`. Data sourced from Economics: National Ocean Watch (ENOW) for the U.S. Territories.

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. L. M. Olea and M. Plagborg-Møller. Local projection inference is simpler and more robust than you think. *Econometrica*, 89:1789–1823, 2021. doi: 10.3982/ecta18756.

J. Olson. Understanding and contextualizing social impacts from the privatization of fisheries: An overview. *Ocean & Coastal Management*, 54(5):353–363, 2011.

C. R. Peña, A. L. M. Serrano, P. A. P. de Britto, V. R. Franco, P. Guarnieri, and K. M. Thomé. Environmental preservation costs and eco-efficiency in Amazonian agriculture: Application of hyperbolic distance functions. *Journal of Cleaner Production*, 197:699–707, 10 2018. ISSN 09596526. doi: 10.1016/j.jclepro.2018.06.227.

J. Piger and T. Stockwell. Differences from differencing: Should local projections with observed shocks be estimated in levels or differences? *Journal of Applied Econometrics*, 40:759–787, 11 2025. doi: 10.1002/jae.70003.

J. M. Piger and T. Stockwell. Differences from differencing: Should local projections with observed shocks be estimated in levels or differences? *SSRN Electronic Journal*, 6 2023. doi: 10.2139/SSRN.4530799.

M. L. Pinsky, G. Reygondeau, R. Caddell, J. Palacios-Abrantes, J. Spijkers, and W. W. L. Cheung. Preparing ocean governance for species on the move. *Science*, 360(6394):1189–1191, 6 2018.

V. A. Ramey and S. Zubairy. Government spending multipliers in good times and in bad: Evidence from US historical data. *Journal of Political Economy*, 126:850–901, 4 2018. doi: 10.1086/696277.

M. N. Reimer, J. K. Abbott, and A. C. Haynie. Empirical models of fisheries production: Conflating technology with incentives? *Marine Resource Economics*, 32:169–190, 4 2017. ISSN 07381360. doi: 10.1086/690677.

A. M. Ricart, J. B. Gómez, R. H. Karm, J. L. Largier, V. B. C. D. Souza, A. S. Dias, M. G. Velázquez, T. Nelson, K. C. Cavanaugh, K. C. Cavanaugh, and B. B. Hughes. Persistent kelp forests during a massive decline reveal the importance of land–sea connectivity. *Ecology*, 106, 9 2025. ISSN 19399170. doi: 10.1002/ecy.70212.

L. A. Rogers, R. Griffin, T. Young, E. Fuller, K. S. Martin, and M. L. Pinsky. Shifting habitats expose fishing communities to risk under climate change. *Nature Climate Change*, 9:512–516, 7 2019. ISSN 17586798. doi: 10.1038/s41558-019-0503-z.

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.

C. D. Romer and D. H. Romer. New evidence on the aftermath of financial crises in advanced countries. *American Economic Review*, 107:3072–3118, 10 2017. doi: 10.1257/aer.20150320.

D. Salgueiro-Otero, M. L. Barnes, and E. Ojea. Climate adaptation pathways and the role of social-ecological networks in small-scale fisheries. *Scientific Reports*, 12, 12 2022. ISSN 20452322. doi: 10.1038/s41598-022-18668-w.

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.

J. Sauer and U. Latacz-Lohmann. Investment, technical change and efficiency: Empirical evidence from German dairy production. *European Review of Agricultural Economics*, 42:151–175, 2 2015. ISSN 14643618. doi: 10.1093/erae/jbu015.

S. A. Sethi, M. Reimer, and G. Knapp. Alaskan fishing community revenues and the stabilizing role of fishing portfolios. *Marine Policy*, 48:134–141, 2014. ISSN 0308597X. doi: 10.1016/j.marpol.2014.03.027.

C. K. Seung. Estimating dynamic impacts of the seafood industry in Alaska. *Marine Resource Economics*, 23:87–104, 2008.

C. K. Seung and E. C. Waters. Measuring the economic linkage of alaska fisheries: A supply-driven social accounting matrix (sdsam) approach. *Fisheries Research*, 97:17–23, 4 2009. ISSN 01657836. doi: 10.1016/j.fishres.2008.12.013.

C. K. Seung and E. C. Waters. Evaluating supply-side and demand-side shocks for fisheries: A computable general equilibrium (CGE) model for Alaska. *Economic Systems Research*, 22:87–109, 3 2010. doi: 10.1080/09535311003677636.

D. A. Smale, M. T. Burrows, A. J. Evans, N. King, M. D. Sayer, A. L. Yunnie, 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.

J. H. Stock and M. W. Watson. Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal*, 128:917–948, 5 2018. doi: 10.1111/ECOJ.12593.

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.

J. J. Van Dijk. Robustness of econometrically estimated local multipliers across different methods and data. *Journal of Regional Science*, 58:281–294, 3 2018. ISSN 14679787. doi: 10.1111/jors.12378.

I. van Putten, C. Cvitanovic, and E. A. Fulton. A changing marine sector in australian coastal communities: An analysis of inter and intra sectoral industry connections and employment. *Ocean & Coastal Management*, 131:1–12, 2016.

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.

H.-J. Wang and P. Schmidt. One-step and two-step estimation of the effects of exogenous variables on technical efficiency levels. *Journal of Productivity Analysis*, 18:129–144, 2002.

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. Technical report, Reef Check Foundation, Marina Del Rey, June 2022. Available at: https://www.reefcheck.org/wp-content/uploads/2022/06/Restoration-of-Northern-Californa-bull-kelp-RCF-final-report-to-OPC.pdf.

B. Watson, M. N. Reimer, M. Guettabi, and A. Haynie. Commercial fisheries and local economies. *Journal of Environmental Economics and Management*, 106:102419, 3 2021. ISSN 0095-0696. doi: 10.1016/J.JEEM.2021.102419.

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.

J. M. Wooldridge. Econometric analysis of cross section and panel data. *MIT Press Books*, 1, 2010. URL https://ideas.repec.org/b/mtp/titles/0262232588.html.

T. Young, E. C. Fuller, M. M. Provost, K. E. Coleman, K. St. Martin, B. J. McCay, and M. L. Pinsky.

Adaptation strategies of coastal fishing communities as species shift poleward. *ICES Journal of Marine Science*, 76(1):93–103, Jan. 2019.

Z. Zhang and J. Ye. Decomposition of environmental total factor productivity growth using hyperbolic distance functions: A panel data analysis for China. *Energy Economics*, 47:87–97, 1 2015. ISSN 01409883. doi: 10.1016/j.eneco.2014.10.023.