Beyond biomass: valuing genetic diversity in natural resource management

Michael R. Springborn* Amanda Faig† Allison Dedrick‡
Marissa L. Baskett§

October 31, 2018

Abstract

Strategies for increasing production of goods from working and natural systems have raised concerns that the diversity of species on which these services depend may be eroding. This loss of natural capital threatens to homogenize global food supplies and compromise the stability of human welfare. We assess the trade-off between artificial augmentation of biomass and degradation of biodiversity underlying a populations' ability to adapt to shocks. Our application involves the augmentation of wild stocks of salmon. Practices in this system have generated warnings that genetic erosion may lead to a loss of the 'portfolio effect' and the value of this loss is not accounted for in decision-making. We construct an integrated bioeconomic model of salmon biomass and genetic diversity. Our results show how practices that homogenize natural systems can still generate positive returns. However, such a system can become only weakly sustainable—the substitution of more physical capital and labor for natural capital must be maintained for gains to persist, weakens the capacity for adaptation should this investment cease, and can cause substantial loss of population wildness. We also introduce an optimization method novel to environmental and resource economics—forward dynamic programming—to solve the model without simplifying restrictions imposed previously.

Key words: biodiversity; portfolio effect; genetic erosion; homogenization; quantitative genetic-bioeconomic; forward dynamic programming; dynamic optimization

*Department of Environmental Science & Policy, University of California Davis, mspringborn@ucdavis.edu
†Aquatic and Fisheries Sciences, University of Washington, afaig@ucdavis.edu
‡Graduate Group of Ecology, University of California Davis, agdedrick@ucdavis.edu
§Department of Environmental Science & Policy, University of California Davis, mlbaskett@ucdavis.edu

1 Introduction

In the production of goods from renewable resources, the boundary between working and natural systems is diffuse. For example, agricultural production on working landscapes takes advantage of natural services like pollination, water regulation and feedstocks. Nominally natural systems such as fisheries and forests are often artificially augmented through hatcheries and tree planting. Natural resource managers have historically made decisions with an eye towards the production of harvestable biomass, i.e. the total mass of crop, timber or fish available. However, in this vein there is growing concern that an emphasis on biomass has come at the cost of homogenizing the species on which these services depend. Timber stands become monocultures (Kelty, 2005). Agriculture suffers genetic erosion of plant and animal species (UNEP, 2005). Hatchery and fishing practices degrade the diversity of stocks (Allendorf et al., 2008). As a result, these systems may be poorly positioned to adapt to a variable environment, e.g. through changes in disease dynamics, climate and availability of food sources. Such systems are at risk of becoming only weakly sustainable (i.e. dependent on man-made capital and labor inputs) or not sustainable at all. An important element of natural capital—genetic variation—has been reduced, driving "homogeneity in global food supplies" (Khoury et al., 2014) and threatening "health, culture and livelihoods" (UNEP, 2007).

In this paper we assess the trade-off between artificial augmentation of biomass in the short-run and the long-run potential for resource homogenization, i.e. degradation of biodiversity underlying a populations' ability to adapt to shocks. Our application involves the augmentation of wild stocks of salmon. Artificial augmentation practices in this system have generated warnings that genetic diversity may be eroding, leading to a loss of the 'portfolio effect' or buffering provided by a diverse set of individuals, and moreover, the value of this loss is not accounted for in decision-making (Carlson and Satterthwaite, 2011). To capture this trade-off, we construct a bioeconomic model of the biomass or quantity of the stock as well as its genetic diversity.

Protection of biodiversity is a cornerstone of conservation and renewable resource management. While the economics literature has historically focused on the problem of maintaining a diverse *set* of species (e.g., Weitzman, 1998) recent research also highlights the importance of diversity *within* a species population (Eikeset et al., 2013; Jardine and Sanchirico, 2015; Zimmermann and Jørgensen, 2015). Genetic differences are the foundation for such diversity, combining with environmental variation to produce a portfolio of traits across individuals. Brock and Xepapadeas (2003) argue that this gene pool of a species is an essential element of natural capital. Although a bioeconomic literature on the evolutionary effects of management has recently emerged, complexity in genetic dynamics has led to compromises in the evaluation of efficient management strategies.

In this paper we implement a dynamic optimization method novel to environmental and resource economics—forward dynamic programming—to solve a quantitative genetic-bioeconomic model without the limitations of earlier analyses. With a few recent exceptions, the representation of genetics in the economics literature is strongly stylized, such that the expression of a

trait is determined by a single gene (or locus) with a small set of alternative types (or alleles). This restricts the trait of concern to be amongst a small, discrete, set of alternatives. For example, to ensure analytical tractability, Brock and Xepapadeas (2003) allow for three pest types while Guttormsen et al. (2008) assume two discrete fish types. Just as a two-period dynamic model is useful for simple insights but insufficient for realistically balancing long-run tradeoffs, so too is a richer model of the genetic continuum needed for representative results.

Continuous-trait genetic model are more flexible and realistic than discrete-trait representations. Consider human height as a canonical example. In a discrete model individuals might be only 'tall' or 'short', while in a continuous model individuals would vary from 'tall' to 'short'. Even though the average is 'medium' in both cases, the population distributions are fundamentally different (and in the discrete case highly unrealistic). Furthermore, the continuous model readily allows for the environment to realistically influence expression of the genes (e.g. nutrition differences driving small to large shifts in observed height).

Jardine and Sanchirico (2015) use such a continuous-trait model to examine how markets incentivize the degradation of early-returning runs of a fishery. While they provide a first look at how economic factors can influence the portfolio of a stock complex, the behavior (and thus genetics) of each subpopulation is assumed to be *fixed*—the only diversity dimension considered is the relative biomass of each subpopulation. Eikeset et al. (2013) and Zimmermann and Jørgensen (2015) relax this constraint in their applications (as we do in our analysis), capturing changes in time-varying traits via a dynamic quantitative genetic model. Both of these analyses identify optimal management given dynamic feedbacks with genetics underlying key traits.

Eikeset et al. and Zimmermann and Jørgensen allow for rich variation in the genetic state of their populations. However, due to strong computational challenges, neither analysis identifies the fully optimal strategy—policies are assumed to follow specific, simple functional forms dependent on two parameters which are optimized through a combination of simulation and brute force search. Our methodology brings the dynamic quantitative genetic-bioeconomic problem back into a dynamic programming framework in which no assumed structure is imposed on the policy function mapping the multi-dimensional stock complex state into management action.

In our application we consider the California Central Valley Chinook (CVC) salmon stock complex. In spring 2008, state and federal fishery managers imposed an emergency closure of Chinook salmon fishing off the coasts of California and southern Oregon due to anticipated poor returns of fall-run CVC salmon. The closure was the first in the fishery's 157 year history and led to an estimated loss of $255 million (Schwarzenegger, 2008). It was argued that at the time CVC salmon lacked the capacity to buffer against unfavorable environmental variation in ocean conditions (affecting food availability) due to a reduction in life history variation (Lindley et al., 2009).

Our key trait of interest—and focus of the genetic model—is migration timing, which is critical for survival via resulting access to food. Each year, young CVC salmon migrate

from upstream spawning grounds, out to the ocean. Survival in the ocean depends on when the young arrive relative to the annual nutrient upwelling near the coast, which attracts an abundance of other species that the young must feed on to grow rapidly (Petrosky and Schaller, 2010; Wells et al., 2012). Young that arrive early face limited food sources and those that arrive late encounter a growing number of gathering predators. While this annual nutrient upwelling typically occurs in the spring and early summer, the exact date varies year to year (Largier et al., 1993). Variation in migration timing is hypothesized to provide "value in avoiding boom and bust dynamics", which would occur if all subpopulations behaved identically, all surviving or succumbing as young (Hilborn et al., 2003).

With respect to management, we focus on a practice that artificially augments biomass, but does so at the risk of degrading genetic diversity. Hatcheries in this system are motivated by the need to compensate for habitat either lost (e.g. due to dams) or degraded. After young salmon are reared in captivity, survivorship is augmented by trucking the young for a portion of their route to the ocean to reduce migration mortality. However, this artificial augmentation practice removes the opportunity for young to imprint on the water chemistry of their home habitat using olfactory cues. The result, when adults return from the ocean to spawn, is increased straying, i.e. individuals diverting from their home stream. Ecologists argue that this homogenizes subpopulations *across* streams (Lindley et al., 2009)[1]. We focus on the ideal artificial augmentation policy, operationalized as the share of the natural migration distance bypassed by hatchery stock. This key management decision has been debated for several years, with assertive arguments being made for limited (or no) trucking due to the possible loss of the portfolio effect (California Hatchery Scientific Review Group, 2012).

This form of augmentation is specific to the system but captures a general tradeoff, tuning a dial that increases current biomass at the potential cost of degrading diversity in the stock portfolio. This belongs to a class of management actions that affect both biomass and genetic diversity, such as in captive breeding, monocropping, and harvest driven fisheries-induced evolution. Our results contribute to a broader understanding of efficient joint management of biomass and biodiversity. These problems highlight the fact that ecosystem management to achieve better short run provision of services can involve key tradeoffs with the longer run goal of ensuing the whole system's capacity to deliver services (Brock and Xepapadeas, 2003).

Incorporating genetic richness across two subpopulations while solving for efficient management policy presents a substantial challenge for standard numerical solution techniques such as value function iteration. We use a recently developed stochastic simulation solution technique that is tractable even under high dimensionality known as approximate dynamic programming or, more intuitively, forward dynamic programming (FDP). This optimization approach first appeared in the operations research and engineering literatures (Powell, 2007; Bertsekas, 2011). To date in economics, the FDP method has been used to solve high-dimension versions of macroeconomic general equilibrium problems by Judd et al. (2011) and Hull (2015).

[1] Across-stream correlations in returns in CVC salmon are larger compared to other salmon complexes and to the era before trucking (Carlson and Satterthwaite, 2011)

In contrast to value function iteration where the solution to the management problem is gradually improved through backwards induction, FDP uses repeated Monte Carlo simulation forward in time to accomplish the same. This method generates much smaller demands on computer memory, allowing solution of high-dimensional problems like the one considered here. Additionally, FDP can circumvent the need for numerical integration and indeed can be used to approximate integration problems that may otherwise be infeasible to solve (Rust, 1997). FDP allows for the application of dynamic programming to "large and complex problems" (Bertsekas, 2011) and can be simple to program relative to the alternatives (Judd et al., 2011).

Our results show that the portfolio is indeed degraded by artificial augmentation—diversity declines within and between subpopulations. Even so, the maximum feasible augmentation is optimal in most cases. The lost portfolio effect value is typically more than compensated for by the population boost from increased survivorship. Further, while we find that a loss in portfolio effect increases the variation in harvest profit it does so only in the positive direction. Despite these economic gains, if artificial production of young ever ceases, having taken advantage of augmentation (trucking) leaves the system in a worse position to recover than if no augmentation had been used. Augmentation also strongly erodes the wildness of both subpopulations. Finally, impacts both good (additional stock) and bad (genetic erosion, loss of wildness) are most pronounced in the subpopulation *without* the hatchery.

Generally our results show how practices that homogenize natural systems through genetic erosion can still generate positive net economic returns. This occurs when the direct demographic boost outweighs the impact of genetic erosion. However, the system is only weakly sustainable—the substitution of more physical capital and labor for natural capital (1) must be maintained for gains to persist, (2) weakens the capacity for adaptation should this capital be removed, and (3) can inadvertently move populations far from their predominantly wild baseline.

In the next section we describe the methods, including the system dynamics, the manager's problem, and the solution method. In the third section we present the results: the optimal value function, the optimal policy function, and simulations to depict how optimal policy compares to alternative policy.

2 Methods

In this section we describe the dynamics of the system, followed by the decision problem, and finally the FDP solution method.

2.1 System dynamics

To describe system dynamics, we follow the biological model and parameterization of Dedrick and Baskett (2018), except where noted. Values used for parameters are summarized in Appendix A.6. We model two genetically distinct subpopulations of salmon in separate streams of the same river system. Let t represent the time period and let salmon subpopulations be indexed by $i \in \{1, 2, 1^*\}$ for the two streams {1,2} and for the hatchery fish $\{1^*\}$ (which are sourced from individuals in stream 1). For each subpopulation, we follow the population size, $N_{i,t}$, and the genetic distribution for migration timing (the key evolving trait) as described by the genetic mean, $\mu_{i,t}$, and variance, $G_{i,t}$. The distribution of observed traits (or phenotypes) in the population is continuous, based on the genetic distribution as well as (fixed) environmental variance. At the individual level, this trait is the migration timing (which can be equivalently expressed as the ocean arrival day given a fixed time to complete the journey). This trait is central to the survivorship of this population as it affects food availability for initial development in the ocean. Figure 1 depicts the model dynamics.[2]

At the beginning of the period (horizontal dashed bar) salmon exit the ocean and swim upstream in order to reach spawning grounds in one of two tributary streams. Both streams have wild spawning grounds. However, we assume that only stream 1 has a hatchery in order to capture feedbacks between subpopulations with artificial reproduction and those without. In order to produce juveniles, the hatchery takes a fixed number of spawners returning to stream 1 (H_{max}) unless the population is low, in which case only a fraction (ζ) of returners are taken[3]:

$$N_{1^*,t} = \min(\zeta N_{1,t}, H_{max}), \tag{1}$$

where $N_{i,t}$ is the number of adult spawners in a subpopulation. The value selected for ζ and all other parameters in the model are listed in Appendix A.6.

We assume wild reproduction is subject to Beverton-Holt density-dependent recruitment, as per Honea et al. (2009):

$$N'_{i,t} = \frac{(N_{i,t} - \mathbb{1}_{i=1} N_{1^*,t}) R_i}{1 + (N_{i,t} - \mathbb{1}_{i=1} N_{1^*,t}) R_i / K_i}, \tag{2}$$

where $N'_{i,t}$ is the number of juveniles produced, $\mathbb{1}_{i=1}$ is an indicator variable for stream 1 (to account for hatchery removals), R_i is the fecundity of wild fish and K_i is the carrying capacity. Reproduction in the hatchery, in contrast, is assumed to occur without density

[2]CVC salmon exhibit overlapping generations with freshwater recruits typically spending two years in the ocean before returning to spawn. However, because the number of state variables is already elevated to capture both biomass and biodiversity, we focus on a single generation to capture tradeoffs as parsimoniously as possible.

[3]This is a simplifying assumption based on expert opinion.

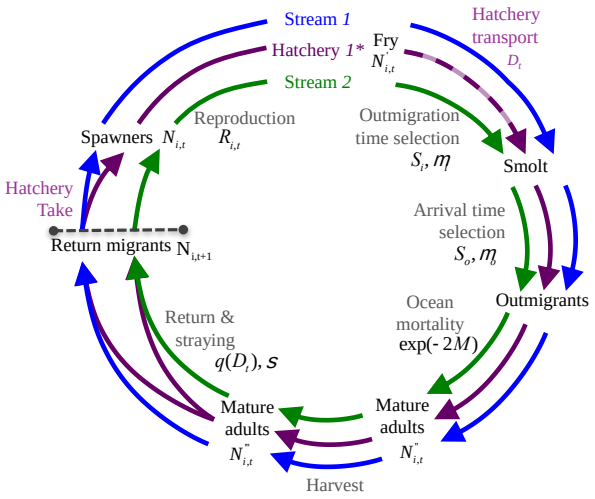

Figure 1: Outline of the model dynamics. Each circle indicates the population rearing environment, $i \in \{1, 2, 1^*\}$. Text along the circle indicates different life history stages (with a dashed bar at the census point), inner text indicates evolutionary and ecological dynamics, and outer text indicates management-driven dynamics.

dependence[4], such that

$$N'_{1^*,t} = R_{1^*} N_{1^*,t}. \tag{3}$$

Each stream contains a genetically distinct subpopulation. Genetic dynamic equations are developed in detail in Appendix A.1. Variation in the key trait—migration timing—across subpopulations is expected because they adapt to local conditions and there is low natural exchange between them.[5] We assume that the genetic mean value of this trait for juveniles is identical to that of their parents. The genetic variance for juveniles is a function of the spawning population variance, assuming random mating of the parents and constant genetic variance at inheritance. After maturing for a time in the stream, juveniles out-migrate towards the ocean. This journey involves migration mortality in the stream that is both trait-independent (e.g. predation) and trait-dependent (i.e. selective) given optimal-timing factors such as stream flow, temperature, etc. Both types of mortality are assumed to depend on the distance the out-migrating juveniles must swim. Wild fish, which are never trucked, have a maximum survivorship of $\kappa_{wild} \in [0, 1]$ (due to trait-independent mortality). Hatchery fish survivorship is "augmented" by bypassing some of the wild journey to the ocean (via trucking) given by the choice variable $D_t \in [0, 1]$, normalized to the unit interval.

[4]This is based on expert opinion to mirror the fact that hatcheries choose the number of adults to take based on capacity and so the hatchery juvenile population size is not set by density-dependence, unlike the wild juvenile population.

[5]Migration timing refers to the date on which the juvenile salmon begins their migration downstream to the ocean. This timing largely determines when the juvenile salmon will arrive to the ocean. In this paper we assume that any difference in migration timing results in an identical difference in ocean arrival timing.

The maximum possible in-stream survivorship of migrating juveniles is thus given by

$$\tilde{\kappa}_i(D_t) = \begin{cases} \kappa_{wild} & \text{if } i = 1, 2 \\ D_t + (1 - D_t)\kappa_{wild} & \text{if } i = 1^*. \end{cases} \tag{4}$$

When survival of hatchery fish is not augmented at all, in-stream survivorship of hatchery fish is identical to that of wild fish (κ_{wild}). If the hatchery augments juvenile survivorship as much as possible ($D_t = 1$), survivorship of hatchery fish at this stage is 100%.[6] Migrating hatchery fish thus have a maximum survivorship that is a convex combination of wild survivorship and full survivorship as a function of augmentation distance (D_t). This migration mortality is non-selective (trait-independent), i.e. affects all fish equally. In order for genetically distinct subpopulations to emerge, there must also be a source of mortality that is selective.

In our model, selective migration mortality is induced by any mismatch between the individual's trait and the ideal trait determined by the ecology of each particular stream, i.e. between each fish's migration date and the optimal migration date for that stream. Each stream has a different ideal outmigration date, due to stream-specific characteristics.[7] Further, ocean selection mortality occurs due to any mismatch between the individual individuals's ocean arrival date (which is determined by their migration date) and the stochastic ocean upwelling date. Finally, there is non-selective natural ocean mortality. $\kappa_{i,t}\left(D_t|\mu'_{i,t}, G'_{i,t}\right)$ represents combined migration and selective ocean survivorship (before any ocean harvest) and is a function of the decision variable and the genotype mean and variance at this time step. See Appendix A.1 for the genetic dynamic equations in full, including how in-stream and ocean selective mortality affect the mean genetic value of the population.

A share of the population, $\exp(-2M)$, survive non-selective ocean mortality for two years before returning (where M is the instantaneous rate of natural mortality). The number of fish that survive and return to spawn as mature adults (in the absence of harvest) is given by:

$$N''_{i,t} = N'_{i,t} \cdot \kappa_{i,t}\left(D_t|\mu'_{i,t}, G'_{i,t}\right) \cdot \exp(-2M). \tag{5}$$

During their time in the ocean, juveniles mature to adults and are subject to harvest at proportion F_t. The number of surviving adults is $N'''_{i,t} = N''_{i,t}(1 - F_t)$.[8] We assume that if the number of aggregate survivors, $N''_t = \sum_i N''_{i,t}$, falls below a "quasi-extinction threshold", $\underline{N}$,

[6] We assume there is no mortality of hatchery fish as they are trucked. This is a simplifying assumption, based on expert opinion and hatchery reports which state, for example, "fish looked very healthy upon arrival and acclimated extremely well" (FFC, 2009). While some mortality is likely we could not find data, or even estimates, as to how much mortality occurs during trucking.

[7] Multiple physical and biological factors can drive this difference. For example, due to differences in air temperature, altitude, and even soil composition between streams (and their watersheds) peak snow melt discharge can vary by weeks between streams (Peterson et al., 2005).

[8] While CVC salmon are typically in the ocean for two harvest cycles, during the first harvest period they are exposed to they are generally under commercial and recreational size limits (O'Farrell et al., 2013, p. 6). Thus we include a single season of mortality from harvest.

the aggregate population goes extinct.[9]

After harvest the remaining adults return to the streams to procreate. Typically they head to the same spawning grounds they were reared in by following scent paths (chemical cues) imprinted while out-migrating (Keefer et al., 2006). Wild salmon naturally stray from their home stream to an alternative stream at some very small proportion, σ. We assume that hatchery salmon stray at a proportion $q(D_t)$ that increases linearly from σ the further they were trucked as juveniles, as per California Hatchery Scientific Review Group (2012):

$$q(D_t) = \sigma + (q_{max} - \sigma)D_t. \tag{6}$$

Returning adults separate as they swim upstream towards the spawning ground, creating populations of size

$$\begin{aligned} N_{1,t+1} &= (1-\sigma)N'''_{1,t} + \sigma N'''_{2,t} + (1-q(D_t))N'''_{1^*,t} \\ N_{2,t+1} &= \sigma N'''_{1,t} + (1-\sigma)N'''_{2,t} + q(D_t)N'''_{1^*,t}. \end{aligned} \tag{7}$$

2.2 The Manager's Problem

Given the complexity of salmon life history across fresh water and ocean environments, there are a number of management decisions that affect their dynamics. Here we focus on the manager's choice of artificial augmentation of juvenile survival, D_t, since this choice is believed to play a central role in both increasing biomass and the persistence (or loss) of genetic diversity between subpopulations.[10] Garnache (2015) found that the gains to improving habitat in a salmon system (through floodplain management) were larger than those from improvements to the fishery management regime.[11] This suggests that efforts to bolster stocks to compensate for habitat degradation are likely to be important decision variables in their own right.

Welfare in the model stems from harvest profits. To specify the harvest exploitation rule, we start with the Pacific Fishery Management Council's plan for Sacramento River Fall Chinook (SRFC) as summarized by Winship et al. (2015). We make two adjustments to this rule in our model. First, we scale the stock level to be consistent with the share of the system we are

[9] This quasi-extinction threshold is based on research by Lindley et al. (2007) who argue that extinction risk is high when the stock falls below a given threshold. For a full description of how this threshold was chosen, see Appendix A.2.

[10] Another potential hatchery decision variable to consider is the quantity of hatchery production. However, since the hatcheries' mandate is generally to produce a roughly constant quantity of juveniles per year, and hatcheries are limited in capacity, we take this level of production as given for this analysis and focus on augmentation (Huber and Carlson, 2015).

[11] This result emerges due to the high marginal benefit of additional habitat, the low opportunity cost of habitat, and the low cost of harvest. Low harvest costs cause the maximum sustainable yield to be very similar to the rent-maximizing total allowable catch, allowing the baseline harvest rate to be near optimum already. The combination of circumstances means the relative gains to habitat improvement are much larger than the gains from harvest policy improvement.

capturing in our focus on two streams: we set the scale of our system at one quarter of the size of the aggregate SRFC stock complex. Second, we smooth the harvest rate function. The rule specified in Winship et al. (2015) is based on a constant escapement approach but includes modifications that introduce several discontinuities. To avoid potentially erratic features in the value and policy functions, we use a smooth approximation to the rule. Further detail on this specification appears in Appendix A.3.

We begin each period as adults make their way upstream after harvest and straying has occurred, but before the hatchery has taken a portion of the returning stream 1 fish. This census point was chosen in order to reduce the number of dimensions the model must track (since at this point, hatchery-reared fish have been absorbed into the subpopulations). Let the state of the system be denoted by $X_t = \{N_{1,t}, N_{2,t}, \mu_{1,t}, \mu_{2,t}, G_{1,t}, G_{2,t}\}$. The augmentation decision is a function of the state at the beginning of the period, $D_t = D(X_t)$.

The stock at the time of harvest is a function of the state, decision and shock: $N_t''(X_t, D_t|\varepsilon_t)$. Given a harvest rate of F, profit is given by:

$$\pi(X_t, D_t|\varepsilon_t) = N_t'' \cdot F(N_t'') \cdot \left(p - c \ln\left(\frac{1}{1 - F(N_t'')}\right)\right), \tag{8}$$

where p is the ex-vessel price per fish, c is a harvest cost parameter and arguments of $N_t''(X_t, D_t|\varepsilon_t)$ have been suppressed on the right hand side for conciseness. Parameterization of the profit function is discussed in detail in Appendix A.4. We assume that p is constant, which is consistent with historical ex-vessel price observations from this fishery. We also make the standard assumption that the marginal cost of harvest is increasing as the stock decreases. Given a discount factor β, the Bellman equation which describes the manager's optimization problem is:

$$V(X_t) = \max_{D_t}\left\{E_{\varepsilon_t}\left[\pi(X_t, D_t|\varepsilon_t) + \beta V(X_{t+1})\right]\right\}, \tag{9}$$

where $X_{t+1}(X_t, D_t|\varepsilon_t)$ is specified in the systems dynamics section above. The optimal management rules are given by optimal augmentation policy function, $D^*(X_t)$.

2.3 The FDP Solution Method

To find the value function of the dynamically optimized system described above, we use relatively new technique known as forward dynamic programming (FDP). We use this solution method because of its ability to handle the many state equations (governing population sizes, genetic means and genetic variances) as well as the stochasticity from the ocean upwelling. Traditional dynamic programming methods (e.g. value function iteration) involve backward iteration that is computationally intensive due to the need to consider the entire state space at each backwards step. In contrast, FDP uses information from simulation of single chains in the state space to iteratively improve the representation of the value function. This

approach is alternatively referred to as neuro-dynamic programming or approximate dynamic programming (ADP) (Hull, 2015). However, we use the terminology FDP since it captures the most striking departure from standard techniques (forward versus backward iterations) and because ADP has been used to refer to other methods (e.g. Simao et al., 2009; Van Roy et al., 1997; Wright and Williams, 1984).

As summarized by Judd et al. (2011), alternative numerical methods (to standard backward iteration) that can solve dynamic stochastic economic models fall broadly into one of three classes: projection methods, perturbation methods, and stochastic simulation methods. FDP belongs to the last of these. All three have their relative advantages and disadvantages and the best choice varies by application. Projection methods, which approximate solutions over a pre-specified domain using deterministic integration, calculate solutions quickly and accurately when there are few state variables, but slow down significantly as the number of state variables increases. Perturbation methods, which find solutions locally using Taylor expansions of optimality conditions, perform well when solving high-dimensional applications, but are limited in the range of their accuracy. Finally, stochastic simulation methods can solve for high-dimensional applications—like the case considered in this manuscript—but are generally less accurate than projection methods and can be numerically unstable. Judd et al. (2011) illustrate how accuracy and stability of the stochastic simulation algorithm can be achieved by normalizing certain variables and modifying the regression step to handle ill-conditioned problems (i.e. through Tikhonov regularization).

We outline our implementation of the FDP solution algorithm in detail below. In broad strokes, the approach involves improving on the current estimate of the value function by forward simulating multiple stochastic chains through the state space and then updating the value function representation by regressing the state vectors on the corresponding "observed" values. The steps to estimate the value function using FDP are as follows:

1. Set FDP parameters.
   (a) Choose the time horizon, T, for each simulation.
   (b) Choose the number of simulations, $\bar{n}$, to complete in a block before executing each regression step and set the regression counter to zero, $z = 0$.
2. Initialize the value function and the state space.
   (a) Set an initial guess for the value function by initializing the parameter vector, θ, of the approximating model, $\bar{V}^{z=0}(X;\theta)$.
   (b) Define a discretization of the state space, $\hat{X}$.
3. For each simulation iteration $n = 1, ..., \bar{n}$, in block $z + 1$ execute the following steps:
   (a) For the initial period, $t = 1$, randomly select a state vector, $X^n_{t=1}$.
   (b) For each period $t = 1, ..., T$, in the simulation execute the following steps:
      i. Compute the next period state vector, $X^n_{t+1}(X^n_t, D_t|\varepsilon_t)$, for simulation n for each possible choice (D_t) and upwelling date (ε_t)
      ii. Calculate the value for the period for every possible choice and upwelling date:

$$v^n_t(X^n_t, D_t|\varepsilon_t) = \pi(X^n_t, D_t|\varepsilon_t) + \beta\bar{V}^z(X^n_{t+1})$$

iii. Find the expected value of each choice by summing a weighted average of the possible upwelling dates:

$$E\left\{v_t^n(X_t^n, D_t|\varepsilon_t)\right\} = \sum v_t^n(X_t^n, D_t|\varepsilon_t) f(\varepsilon_t),$$

where $f(\varepsilon_t)$ is the probability mass function for upwelling date, ε_t.

iv. Choose D_t^* to maximize expected value

$$\max_{D_t} E\left\{v_t^n(X_t^n, D_t|\varepsilon_t)\right\}.$$

v. Randomly select the upwelling date for this period, ε_t^*.

vi. Compute the value of having chosen D_t^* conditional on the actual shock:

$$\tilde{v}_t^n = v_t^n(X_t^n, D_t^*|\varepsilon_t^*)$$

vii. Set the step size, $\delta_t \in [0, 1]$, specifying the weight placed on $\tilde{v}_t^n$ in the updating step that follows.

viii. Compute the expectation with a linear combination of the newly calculated value and the previously approximated value at state X_t^n:

$$\widetilde{V}_t^n = \delta_t \tilde{v}_t^n + (1 - \delta_t)\bar{V}^z(X_t^n).$$

(c) After completing $\bar{n}$ simulations of T periods each, there are $\bar{n} * T$ observations of the state vector visited ($\mathbf{X}$) and the associated updated value estimate ($\widetilde{\mathbf{V}}$). Scale and center the data and regress $\mathbf{X}$ on $\widetilde{\mathbf{V}}$.[12]

(d) Increment the regression counter by one ($z = z + 1$) and define $\bar{V}^z$ as the fitted model from the regression.

(e) Check for convergence. Calculate the maximum relative deviation between the current and former value function estimate: $\Delta_z = \max\{(\bar{V}^z(\hat{X}) - \bar{V}^{z-1}(\hat{X}))/\bar{V}^{z-1}(\hat{X})\}$. Let $\Delta_{\bar{z}}$ represent the average of Δ_z over the last m regressions. If $\Delta_{\bar{z}} < \omega$, the convergence criterion is met and the program can be terminated.[13] Otherwise, repeat step 3.

After convergence, the final optimal policy function $D(X_t)$ is computed using the final estimate of the value function above, $\bar{V} = \bar{V}^z$.

[12] In our application we do not find that the data is ill-conditioned for regression. However, when this is not the case a Tikhonov regularization can be used to facilitate regression as follows. Replace the $\mathbf{X}'\mathbf{X}$ component of the standard OLS equations with $(\mathbf{X}'\mathbf{X} + \eta I_n)$, where η is a very small number (e.g. 10^{-5}) and I_n is the identity matrix, such that $\beta = (\mathbf{X}'\mathbf{X} + \eta I_n)^{-1}\mathbf{X}'\tilde{\mathbf{V}}$. An equivalent approach involves simply augmenting (appending) the matrix ηI_n to the data matrix $\mathbf{X}$, and adding a vector of zeros of the same length to $\tilde{\mathbf{V}}$.

[13] Using the average maximum deviation over several iterations ($m > 1$) helps avoid premature stopping that may result when a pair of regressions happen to produce similar results. We iterate until $\Delta_{\bar{z}} < \omega = 0.25\%$ which results in an average absolute deviation across all nodes considered in the state space of 0.005%

Implementing the FDP algorithm above requires making choices over a set of solution method parameters and functional forms, which are detailed in Appendix A.5. A central challenge in any dynamic programming problem is to implement a representation of the value function using either a lookup table, parametric model or non-parametric model (Powell, 2011, p. 233). Existing applications have used either a lookup table (e.g. Hull, 2015) or a parametric model (e.g. Judd et al., 2011; Maliar and Maliar, 2013). A lookup table for the value function defined at discrete values is simple in that it does not involve assuming any special structure in the value over the state variables. But unless states are naturally discrete at level of coarseness of the lookup table, the solution will require aggregation whereby intermediate levels are collapsed around the neighborhood of a node. This creates bias in the value function, specifically a departure between the true value at a point and the aggregate value for its neighborhood (Powell, 2011, p. 299). In contrast, parametric (e.g. polynomial) models exploit structure in the value over the state variables (Powell, 2011, p. 304). The advantage is that fewer points are needed and optimization is accelerated by the increased smoothness (Judd, 1996). However, as Powell (2011, p. 316) summarizes, the promise of parametric models is countered by a key handicap: "they are only effective if you can design an effective parametric model, and this remains a frustrating art".

To address the weaknesses of both lookup tables and parametric approximations, we implement a nonparametric representation of the value function. This generates a continuous function that allows for very flexible behavior without the need to choose specific basis functions or structure. Powell (2011, p. 316) observes that nonparametric methods hold "tremendous" promise but face substantial challenges. We take advantage of recent advancements in nonparametric statistics (Meyer and Liao, 2016) which allow the "data" to establish both the structure and actual levels of the function.[14] Parametric approximation of the value function can lead to instabilities in the solution procedure from poor interpolation between the nodes, for example because overall shape is not preserved (Judd, 1996) which can result in loss of a convex optimization problem needed to identify a global optimum (Cai et al., 2017). The nonparametric approach allows for shape preservation (concavity and monotonicity) if there is reason to believe such properties will hold. This feature has also been pursued in the realm of parametric models. For example Cai et al. (2017) modify a Chebyshev polynomial approach to impose shape constraints at the nodes, though global shape preservation is not assured and must be tested and nodes tuned. In contrast, the approach used here ensures shape constraints are met and no tuning is needed.

In our model for the value function, each state variable enters additively in nonparametric form. Because we expect the direct contribution of biomass to always be positive, we constrain the direct contribution of N_1 and N_2 to the value function to be increasing. We also include several interaction terms but do so parametrically. Thus, while the core of the value

[14]Specifically, we use the constrained generalized additive model (CGAM) routine of Meyer and Liao (2016) based on results from Meyer (2008). This approach generates a maximum likelihood estimator that is identified using an iteratively reweighted cone projection algorithm. The CGAM routine also allows for the imposition of various constraints over shape (concavity/convexity), monotonicity (increasing/decreasing), and smoothing. These constraints establish a convex cone, which motivates use of the cone projection algorithm. CGAM implements a nonparametric regression using ordinal basis functions serving as the regressors in an ordinary least-squares model with a Gaussian error term.

function model is nonparametric, the set of interactions enters additively and parametrically for computational reasons, explained in detail in Appendix A.5.

3 Results

We present results below in three parts. First, we describe the value function estimate and implications for the value of biodiversity. Second we summarize the policy function characterizing optimal management. Finally, we present simulation results to illustrate expected outcomes in the system under the optimal augmentation policy compared to no augmentation.

3.1 Value Function

In Figure 2 we present the estimated value function over stock levels for each subpopulation with all other states set to their modal levels. As expected the value function generally increases with the population of either stock, since additional biomass is typically better for future harvest profits. The value of an additional individual (in either stream) is high when there is a low population, which results in a steep value function at low population levels. Due to the density dependent nature of recruitment, as the population increases the value function exhibits an "elbow" shape, after which the marginal value of additional fish diminishes abruptly and becomes negligible.[15] This shape illustrates the need for the nonparametric approach we take to model these curves since such an elbow shape is particularly difficult to model with a parametric functional form. It might be tempting to address this challenge with a piecewise polynomial function. However this is cumbersome if not unworkable since the location of the elbow can depend on the level of other states and evolves as the iterative solution technique proceeds.

In Figure 2 we see that the marginal value of reproducing adults in the stock with a hatchery (N_1) is larger (until the marginal value becomes negligible). These reproducing adults not only contribute to natural reproduction but uniquely also contribute to hatchery production.[16]

[15] The young a spawning ground can support is limited by available space and resources. The first few thousand individuals have a direct impact on the population size of the subsequent generation, and in diminishing the likelihood the population will go extinct. Once enough adults return to the spawning ground, additional individuals contribute little to the population size of the subsequent generation. Recall that the census point at which we model the value function is after harvest and straying has taken place but before hatchery take and spawning (see Figure 1).

[16] At high stock levels the value can decreases with additional stock. This effect is very small though nonetheless present. It is not due to the direct effect of each stock on value (which is constrained to be positive) but rather the interaction between the two stocks: the marginal value of stock in one stream is decreasing in the stock level of the other stream. This is shown by the coefficients on the parametric population interaction term from the regression model. As detailed in Appendix A.5, the value function model is nonparametric *except* for the interaction terms, which are parametric. The coefficient on the N_1N_2 interaction term is negative and significantly different from zero (p-value < 0.01).

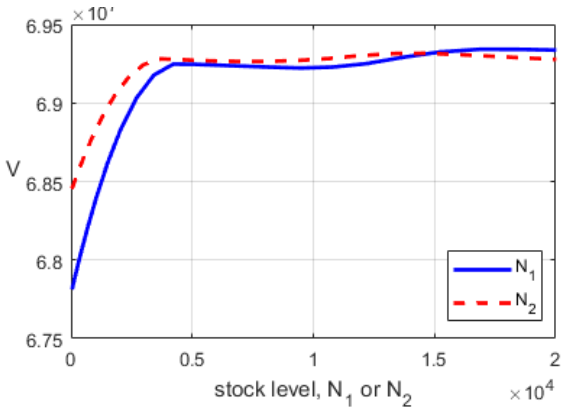

Figure 2: The value function over stock levels (horizontal axis) in each subpopulation (1 and 2) All states that do not vary in a panel are set to their modal levels.

In Figure 3 we present the value function over each population's mean genetic trait (migration day), μ_i. We present these curves for different levels of genetic variance since there is an interaction in the value of mean and variance. Each population has a mean genetic state that maximizes future value, V. The ideal mean trait would be a level that maximizes survivorship given mortality in migration (specific to a subpopulation) and ocean arrival (shared). For subpopulation 2 we see this to be the case: the value function peaks between the trait level best suited for the ocean and migration (see vertical lines). But for subpopulation 1, artificial augmentation means that the ideal is suited almost exclusively for the ocean alone.[17] The broader insight is that artificial augmentation that bypasses elements of the natural life cycle undermines incentives to maintain a population that is adapted for those natural conditions. We show later in our simulation analysis how this leads to genetic erosion that leaves the population in poor shape should managers ever seek a return to fully natural reproduction.

The various curves in Figure 3 provide an initial sense of the value of increasing diversity in the form of genetic variance (i.e. value of higher G_i). We isolate this relationship in Figure 4 where we plot the percentage change in value as the genetic variance of each subpopulation increases. For each subpopulation, we show two cases, one in which the mean trait of the population (μ_i) is ideal for the value of the stock and another in which it is poor. Two results emerge. First, the value of variation in this case is fairly small, leading to a increase in value of 3% at most. Second, when each subpopulation genetic mean is near its ideal level, the value of variation is very small or even slightly negative. Thus when the genetic mean is best adapted to the conditions faced by a subpopulation, variance is of little (and possibly negative) value. This latter result is intuitive but the weak value of variance in general is surprising since it is conventionally believed to be quite valuable, at least when when a population's genetic mean is not already ideally positioned. Qualitatively similar results hold for a range of cases

[17]The value function is typically highest when subpopulation 1 has a mean migration date of day 49.5 and stream 2 fish, 52.7. The only exception is at high levels of G_1 which almost never occur when the system is simulated (see Appendix B.2). A large portion of stock 1 fish do not undergo in-stream selection (day 43 is ideal for migration in stream 1) when they are trucked, whereas all stock 2 fish do experience in-stream selection (day 57 is ideal for migration in stream 2).

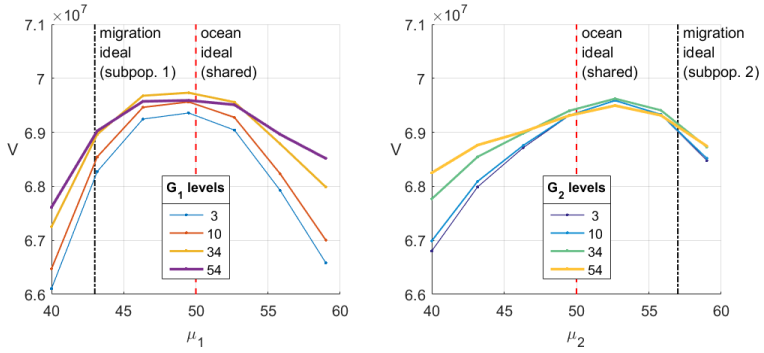

Figure 3: The value function over the mean genetic state (μ_i) subpopulation 1 (left panel) and 2 (right panel). Multiple curves within each panel depict the value at various levels of genetic trait variance (G_i), from largest (thickest line) to smallest (thinnest line). All states that do not vary in a panel are set to their modal levels. Dash-dot vertical lines indicate the ideal genetic state for migration in subpopulation 1 (43) and 2 (57); dashed vertical lines indicate the ideal genetic state for ocean conditions (50), which is shared.

as shown in Supporting Figure 9 (Appendix B.1).[18] In general, we observe that the value of trait diversity is small and depends on the genetic mean—the better adapted the existing population is, the less value there is in variation.

3.2 Policy function

Despite concerns of a degraded portfolio, across most of the state space we find that the optimal policy is to set artificial augmentation at its maximum (trucking hatchery fish over their entire migration distance). For example, across the discretization of the state space used in the FDP solution process, full augmentation is optimal at 82% of the loci. No augmentation is optimal in 7% of the loci, and some interior level at 11%. These percentages will change depending on the particular discretization considered but quickly provide a coarse summary of the preponderance of optimal management choices.

Optimal augmentation that is less than maximum (interior or zero) is more common when the subpopulation with the hatchery (stream 1) is low. This is somewhat surprising since augmentation serves to boost survivorship of these individuals. However it does so at the cost of increased spillover (straying) to the other subpopulation, just when these individuals are needed most. Less than maximum augmentation is also more common when subpopulation 1 is well adapted to face its unique migration mortality.[19] The policy function illustrating this case is presented in Figure 5. For subpopulation 1 (left panel) we see that less augmentation is preferred when μ_1 is lower, i.e. near its ideal for minimizing migration mortality. Intuitively

[18]Further insight into this relationship comes from the coefficients on the parametric interaction terms ($\mu_i G_i$) from the value function regression model. Supporting the graphical analysis above, we find that as variation (G_i) increases, value (V) falls fastest when the mean trait (μ_i) is near its ideal (depicted graphically in Supporting Figure 10, Appendix B.1).

[19]I.e. the genetic mean is near or just below the in-stream 1 ideal, $\mu_1 = 43$.

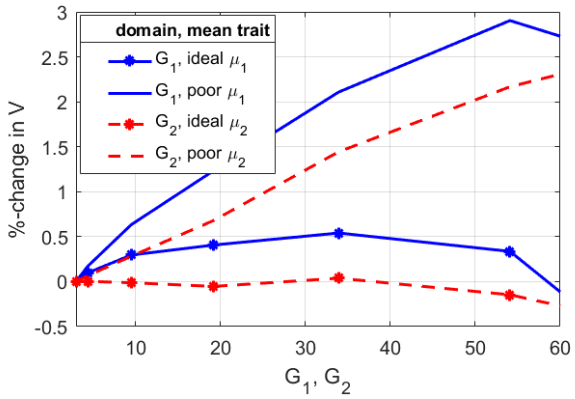

Figure 4: Relative percentage increase in value as a function of the genetic trait variance in the subpopulation (G_1, G_2) for levels of the mean genetic trait state (μ_1, μ_2) that are approximately ideal (49.5, 52.7) or poor (59, 40). All states that do not vary in a panel are set to their modal levels.

this effect is stronger when the population is tightly distributed around this mean (lowest G_1). Here the relative returns to augmentation here are low since natural migration mortality is at its lowest.

For the non-hatchery stream (2), maximum augmentation is optimal more frequently when N_2 is very low. This demographic driver is intuitive since when N_2 is low there is a stronger recovery value to spillovers (straying hatchery fish). In the right panel of Figure 5, we see that more augmentation is appealing when the mean trait (μ_2) is at an ill-suited extreme (high or low)—in this case more strays from stream 1 (induced by augmentation) are useful to bring stream 2's mean trait towards the center. This effect is also strongest when the stream 2 population is tightly distributed around any given mean (i.e. G_2 is low).

Overall optimal management depends on a combination of demographic (biomass) and genetic effects. However, in this case demographic implications dominate, for example, through immediate survivorship of subpopulation 1 migrators and recovery implications for both threatened stocks. This suggests that despite loss from genetic erosion, such degradation may be compensated for by demographic advantages, an outcome demonstrated explicitly below.

3.3 Simulation analysis of dynamic outcomes

To explore implications of the optimal policy we generate Monte Carlo simulations of the system using 3,000 repetitions over 50 periods. For comparison with the optimal policy we also consider the case in which no artificial augmentation is used.[20] While starting points in the state space are chosen randomly, mean paths for the simulated variables stabilize by

[20]Simulation results for a policy of maximum augmentation in all cases (not pictured) are essentially indistinguishable from those under the optimal policy.

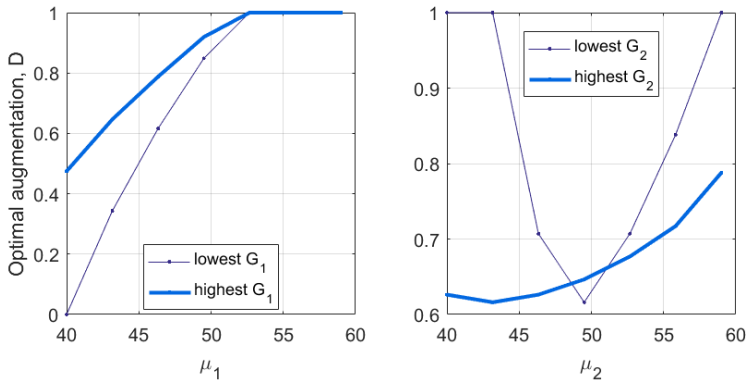

Figure 5: Optimal augmentation policy as a function of μ_i for subpopulation 1 (left) and 2 (right). All hidden state variables are at their modes except for N_1 which is set as the lowest non-zero value in the FDP grid ($N_1 = 42$).

	Change in mean outcome				
	genetic mean, μ	genetic variance, G	population, N	profit, π	wild origin share
subpop. 1	0.1	-1.0	3%	.	-0.063
subpop. 2	-2.3	-1.4	23%	.	-0.382
aggregate	.	-2.8	11%	16%	.

Table 1: Change in the mean outcomes (in percentage or raw terms) due to a shift from none to optimal augmentation. Statistics represent simulated averages over time after excluding the first 30 periods for burn-in.

period 30, as shown in Supporting Figure 11 (Appendix B.2). In the analysis below, we exclude this initial burn-in time frame and report results based on average outcomes over the remaining 20 periods.

In Figure 6 we show simulation results for each state variable for subpopulation 1 (top row) and 2 (middle row) as well as the aggregate population, profit and variance (bottom row). Specifically, we plot the cumulative mass functions (CMFs) for each variable to show how the distribution of outcomes differs between the optimal (solid line) and no augmentation (dashed line) policy. We focus on CMFs since we are concerned with the potential for both expected outcomes and extreme, boom or bust outcomes. The mean outcome level in each case is depicted with a star. In the accompanying Table 1 we summarize the change in these mean outcome levels due to shift from no augmentation to optimal augmentation. Before we consider specific variables, the table highlights a surprising overall pattern that holds for all the key outcomes: although artificial augmentation centers on subpopulation 1, the largest impacts in every case are spillover-driven effects experienced by subpopulation 2. This illustrates that it can be the spillover effects on other populations (both good and bad) that dominate the bottom line, rather than the more obvious direct effects of management on the target population.

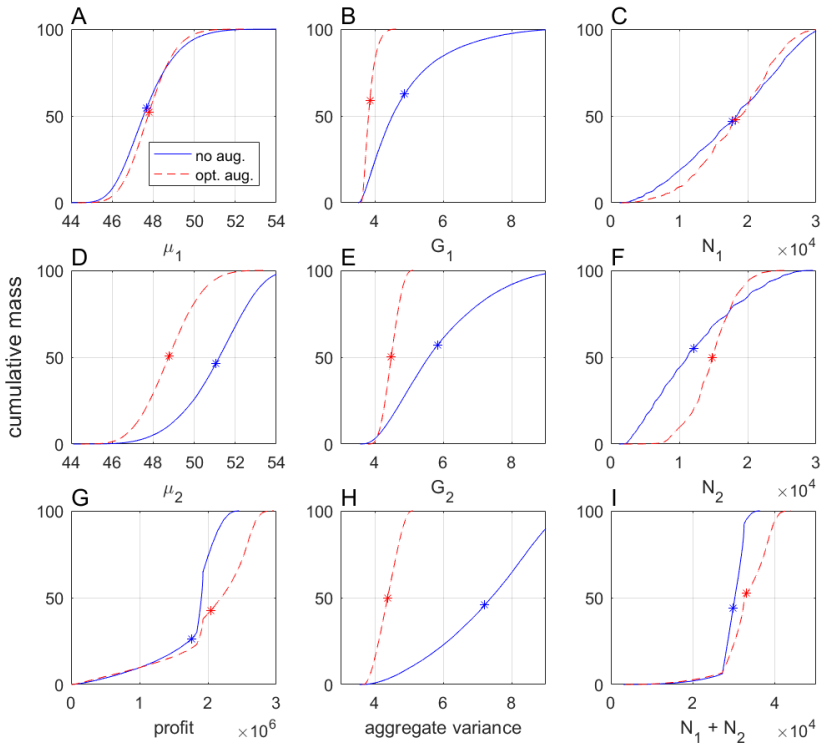

Figure 6: Cumulative mass functions showing likelihood of outcomes for state variables for subpopulation 1 (top row) and 2 (middle row) as well as aggregate variables (bottom row) under no augmentation (solid lines) and optimal augmentation (dashed lines). For the 3,000 simulation runs of 50 periods, the first 30 periods are excluded for burn-in. Stars indicate mean outcome level.

3.3.1 Characterizing portfolio loss

Considering specific variables, we begin by highlighting substantial shifts in the genetic makeup of the stock complex. Panels A and D show that optimal augmentation causes the genetic mean for both subpopulations to converge towards each other (dashed curves in panels A and D are closer together than solid curves). The effect is stronger for subpopulation 2 and is due to increased spillovers from subpopulation 1.[21] This loss of between-population diversity is mirrored by a loss of within-population diversity. In panels B and E we see that the expected genetic variance shrinks for both subpopulations under optimal augmentation (dashed curves shift left). Again, as summarized in Table 1, this effect is stronger in subpopulation 2, which experiences spillover effects (rather than subpopulation 1 which experiences the direct effects of artificial augmentation). We can also calculate a metric of aggregate genetic variance for the stock complex, which is a function of all six state state variables (see Dedrick and Baskett, 2018). We show this measure in panel H—mean aggregate genetic variance falls from 7.2 with no augmentation to 4.4 under optimal augmentation.

Overall we see a clear loss of heterogeneity in the system from artificial augmentation, both in converging genetic means and falling genetic variances. Under the logic of the portfolio effect, we might expect such loss of diversity to increase boom and bust cycles, possibly leading to lower returns in the complex. However, we find that the opposite is the case. We show in panels C and F that optimal augmentation *narrows* the range of likely stock levels for both subpopulations. Both extreme low and extreme high stock levels are less likely. The net effect is an increase in mean stock sizes for both streams. Again, these effects are most pronounced for the for the population *without* the hatchery (2).

3.3.2 Net effects on biomass

We might also expect the aggregate stock outcomes ($N_1 + N_2$) to mirror outcomes in the subpopulations. Instead we see in Figure 6, panel I, that the likelihood of extreme low stocks is essentially unchanged, while the likelihood of extreme high stocks substantially increases. Two dynamics drive this outcome. First, in panels C and F we see that the likelihood of low population levels is dimished more than for high populations, especially for subpopulation 2. Second, the management intervention reduces the pronounced, inverse coupling of the two populations—a strong negative correlation between subpopulation levels under no augmentation (-0.91) weakens (-0.45) under optimal augmentation.[22] No longer is a strong performance in one stock so tightly connected to poor performance in the other.

Stock outcomes suggest that the value of substantial augmentation in terms of additional biomass outweighs the cost of porfolio loss. This is confirmed in panel G where we present

[21] Subpopulation 1's genetic mean increases because artificial augmentation reduces migration selection which favors earlier timing.

[22] The strong counter-movement under no augmentation occurs because when shared environmental conditions (in the ocean) are particularly suited for one subpopulation (e.g. early food availability favoring early arrivers from subpopulation 1) they are ill-suited for the other.

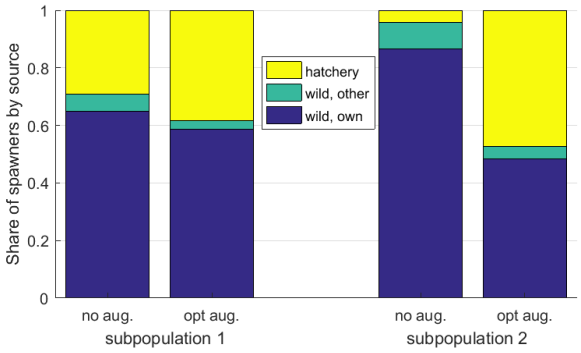

Figure 7: Mean share of spawner source for subpopulation 1 (left) and 2 (right) under no augmentation and optimal augmentation. ("wild"=naturally reared, "own"=from the same subpopulation, "other"=from other subpopulation.)

the CMF for profits.[23] Optimal augmentation raises mean profits by approximately 16%. A remaining concern is that volatility of profits will also rise due to loss of the diversity in the portfolio. We see in panel G that the spread of profit levels is indeed greater with optimal augmentation. However, the additional variation is almost exclusively *upside risk*; while the lower section of the CMF for profit remains essentially fixed, the upper section shifts towards higher values.[24]

Another potential unintended side-effect of hatchery management is loss of wildness, i.e. the replacement of wild-reared individuals with those from the hatchery. In Figure 7 we summarize the origin of reproductive adults. Specifically, we present the share that come from wild-reared juveniles versus hatchery stock (averaged across simulation runs). Relative to no augmentation, the optimal augmentation policy (mostly maximum augmentation) substantially increases reliance on artificial hatchery stock. For the subpopulation 1, hatchery spawners increase from an average of 29% to 38%. More strikingly, for subpopulation 2 (with no hatchery), hatchery spawners increase from an average of just 4% to 47%. In fact, under optimal augmentation with the associated increase in spillovers, the majority of subpopulation 2 spawners are from the other subpopulation—either hatchery-origin (47%) or wild (4%).

3.3.3 Implications of a degraded portfolio

To explore implications of the degraded portfolio if were to cease, we extended each of the 3,000 simulations from above for another 30 periods with no hatchery production (and therefore no augmentation either). In Figure 8 we present mean outcomes for each state variable and for profit. The state variable panels (first three panels) show that without artificial augmentation, the system adjusts to a hatchery shutdown faster and more smoothly: population

[23] Profit is not separated by stream since the ocean fishery combines both populations.

[24] For further perspective into how optimal augmentation changes likely outcomes, in Appendix B.2 we show the relative frequencies for both subpopulations at once for each state variable.

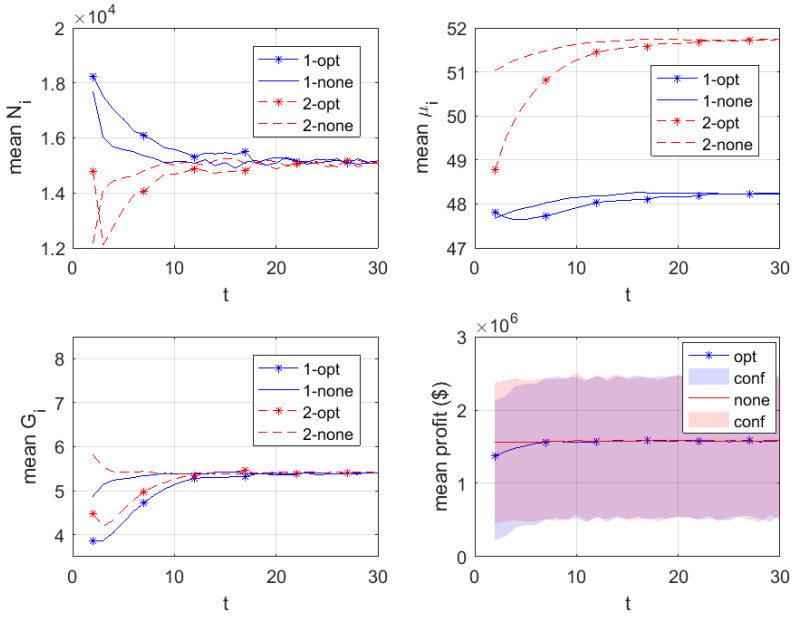

Figure 8: Mean levels for state variables (for subpopulations 1 and 2) and profit once hatchery production ceases, if the hatchery was optimally ('opt') or not ('none') augmentated prior to shutting down, over 3,000 simulation runs across 30 periods. In the final panel, 90% confidence intervals ('conf') are shown around mean profit.

sizes and trait means and variances equilibrate in 10-15 periods (versus 20-25). As discussed in the previous section, augmentation increases stock size while degrading the portfolio. As a result of augmentation (and spillovers), subpopulation 2 is particularly poorly suited to local conditions, which can be seen by the large adjustment in the trait mean (μ_1) that follows the hatchery shutdown (top right panel). As a result, initial survivorship of subpopulation 2 stock is poor (top left panel). When the hatchery was operational it more than compensated for the poor survivorship of wild fish in subpopulation 2. When hatchery production stops, so too does this compensation, resulting in a swift initial decline in subpopulation 2.

The bottom right panel of Figure 8 shows that when the hatchery shuts down, if there was no augmentation prior to closure then annual profit drops directly to its new mean. Alternatively, under pre-closure augmentation, the drop in stock is greater and the expected profit lower, compared to the no augmentation case. Wild stocks do eventually adapt and expected profit recovers after a number of periods. Overall a degraded portfolio leaves the system ill-suited for a return to fully natural production, though it eventually does recover.

Our quantitative genetic model supports wide flexibility in the level of the trait of interest. A limitation of this approach is that it does not capture potential irreversibility in loss of the portfolio. Regardless of the degradation, the model allows eventual recovery to a genetic state that is adapted to a particular scenario. This arguably results in an optimistic picture of capacity for recovery. Our model might overstate the amount of likely genetic diversity available for adaptation to changing conditions. For example, at small or highly variable

population sizes, particular genes may disappear due to stochasticity in survivorship and reproduction (i.e. genetic drift and bottlenecks) (Lande, 1998).

Unfortunately it is not feasible to capture these complexities within our current model (without vastly increasing the dimensionality of the state space). However, as a proxy experiment we repeat the hatchery shutdown simulation experiment with one change: the genetic state in each simulation remains fixed at the level observed immediately before the shutdown. This provides a window into a particularly pessimistic scenario in which the capacity for subsequent adaptation has been lost. Relative to no augmentation, having been optimally augmenting before the hatchery shutdown leads to a lower aggregate mean stock that is also highly skewed towards subpopulation 1 (at a ratio of 2:1) and thus an ongoing mean profit that is 17% lower. This mean penalty for having degraded the portfolio through artificially augmentation is similar in magnitude to the mean profit boost augmentation provides while the hatchery is active. The perfect irreversibility of portfolio degradation in this scenario is an extreme case. However it suggests that if potentially valuable genes are lost—e.g. from augmentation used to increase fishery profits in the short term—some level of productivity would be permanently lost should the hatchery be retired.

4 Discussion

Our results illustrate how practices that homogenize natural resources through loss of genetic diversity can still generate net returns from a profit perspective. On one side of the tradeoff, augmentation erodes diversity both within subpopulations and between them. However, the lost portfolio effect is outweighed by the demographic effect (direct increase in biomass). The resulting configuration is only weakly sustainable: while returns are bolstered by substituting more physical capital and labor for natural capital, additional dependencies and unintended consequences result. The physical capital must be maintained for gains to persist, the capacity for adaptation should this capital be removed is weakened, and wildness of populations can dramatically fall.

A variety of U.S. policies express a societal value for species wildness in the form of natural genetic diversity. For example, under the Endangered Species Act, each genetically distinct and reproductively isolated salmon stock (evolutionarily significant unit) receives separate consideration for listing due to its unique contribution to the "evolutionary legacy" of the species (Waples, 1991). In addition, hatchery management plans for California's fall-run Chinook salmon have the mitigation of "impacts to naturally produced salmonids" among their goals, including strategies to "reduce ecological and genetic interactions" between hatchery and wild fish (e.g., Lee and Chilton, 2007). While loss of wildness is qualitatively understood to be problematic, operationalizing this loss of natural capital such that the ecosystem service impact can be endogenized in the bioeconomic model is a promising track for future research.

It is important to acknowledge that several assumptions made could lead to understatement of the genetic consequences of augmentation. First, in such systems particular genes may

completely disappear due to stochasticity in survivorship and reproduction. This leads to a loss of overall genetic diversity and therefore adaptive capacity (Lande, 1998). In certain circumstances the model used here could overstate the amount of genetic variation remaining in the population. Accounting for the resulting genetic drift and bottlenecks of gene loss would require following many individual genes rather than the overall genetic mean and variance as we do here. Second, hatcheries impose domestication selection on a variety of traits that can reduce survivorship and reproductive success in wild populations when interbreeding occurs (Araki et al., 2007; Reisenbichler and Rubin, 1999). Accounting for additional fitness consequences of hatcheries would require following multiple co-evolving traits. By focusing on a single trait, as we do here, might underestimate the demographic and productivity consequences of hatchery-wild interbreeding. Finally, we focus on two subpopulations; having a more realistic representation of a larger number of streams would increase the amount of genetic variation across streams that hatchery practices could affect.

A number of additional assumptions made for tractability bare discussion as caveats and avenues for further research. First, our environmental shock is stationary. One might expect the value of diversity to be higher in the presence of regime shifts that change favored behavior. We did test to see how optimal management would change under non-stationary, Pacific decadal oscillation (PDO) driven shocks. However, we find that the qualitative results are similar despite the possibility of large intermittent shifts in ideal behavior. Also with respect to the shock process, it may be that an important source of value to adaptive capacity stems from directional environmental change (i.e. shifts in a consistent direction) such as under climate change, which was not explored here.

We also treat harvest as a separate and fixed management variable; our harvest control rule is a stylized version of the current rule used in this system. Previous work in the same system (Garnache, 2015) has shown that returns to optimizing management in *non-harvest* elements of the life cycle depend on whether harvest effort is rationalized. Future work will determine how important harvest is in managing the portfolio effect, and how an optimal harvest policy affects augmentation policy. Other likely relevant decision variables include the timing of water releases from dams and the output level of the hatchery.

Finally, this paper also introduces the use of forward dynamic programming (FDP) to solve bioeconomic problems of substantial dimension. FDP facilitates, for the first time, dynamic optimization of a model with a continuous genetic trait without any simplifying constraints on the policy function. Comprehensive management of natural resources means accounting for multiple dimensions of the natural capital stock, including multiple stocks of biomass, genetic diversity and information. FDP presents a useful new tool for expanding the set of feasible problems that can be solved.

References

Allendorf, F. W., P. R. England, G. Luikart, P. A. Ritchie, and N. Ryman (2008). Genetic effects of harvest on wild animal populations. *Trends in Ecology and Evolution 23*(6), 327–337.

Araki, H., B. Cooper, and M. S. Blouin (2007). Genetic effects of captive breeding cause a rapid, cumulative fitness decline in the wild. *Science 318*(5847), 100–103.

Bertsekas, D. P. (2011). *Dynamic programming and optimal control 3rd edition, volume II*. MIT, available at http://web.mit.edu/dimitrib/www/dpchapter.pdf.

Brock, W. and A. Xepapadeas (2003). Valuing biodiversity from an economic perspective: A unified economic, ecological and genetic approach. *American Economic Review 93*(5), 1597–1614.

Cai, Y., K. L. Judd, T. S. Lontzek, V. Michelangeli, and C.-L. Su (2017). A nonlinear programming method for dynamic programming. *Macroeconomic Dynamics 21*(2), 336–361.

California Hatchery Scientific Review Group (2012). California hatchery review report. Technical report, Prepared for the US Fish and Wildlife Service and Pacific States Marine Fisheries Commission.

Carlson, S. M. and W. H. Satterthwaite (2011). Weakened portfolio effect in a collapsed salmon population complex. *Can. J. Fish. Aquat. Sci. 68*(9), 1579–1589.

Dedrick, A. G. and M. L. Baskett (2018). Changing connectivity and population stability: quantifying the effects of hatchery trucking on the portfolio effect in salmon.

Eikeset, A. M., A. P. Richter, E. S. Dunlop, U. Dieckmann, and N. C. Stenseth (2013). The economic repercussions of fisheries-induced evolution. *PNAS 110*(30), 12259–12264.

Fishery Foundation of California (FFC) (2009). Final Report: San Francisco Bay estuary acclimation of Central Valley hatchery raised Chinook salmon project. Technical report, FFC, Elk Grove, CA.

Garnache, C. (2015). Fish, farmers, and floods: Coordinating institutions to optimize the provision of ecosystem services. *Journal of the Association of Environmental and Resource Economists 2*(3), 367–399.

Guttormsen, A. G., D. Kristofersson, and E. Nævdal (2008). Optimal management of renewable resources with darwinian selection induced by harvesting. *Journal of Environmental Economics and Management 56*, 167–179.

Hilborn, R., T. P. Quinn, D. E. Schindler, and D. E. Rogers (2003). Biocomplexity and fisheries sustainability. *Proceedings of the National Academy of Sciences of the United States of America 100*(11), 6564–6568.

Honea, J. M., J. C. Jorgensen, M. M. McClure, T. D. Cooney, K. Engie, D. M. Holzer, and R. Hilborn (2009). Evaluating habitat effects on population status: influence of habitat restoration on spring-run chinook salmon. *Freshwater Biology 54*(7), 1576–1592.

Huber, E. and S. M. Carlson (2015). Temporal trends in hatchery releases of fall-run chinook salmon in california's central valley. *San Francisco Estuary and Watershed Science 13*(2), article 3.

Hull, I. (2015). Approximate dynamic programming with post-decision states as a solution method for dynamic economic models. *Journal of Economic Dynamics and Control 55*, 57–70.

Jardine, S. L. and J. N. Sanchirico (2015). Fishermen, markets, and population diversity. *Journal of Environmental Economics and Management 74*, 37–54.

Judd, K. L. (1996). Approximation, perturbation, and projection methods in economic analysis. *Handbook of computational economics 1*, 509–585.

Judd, K. L., L. Maliar, and S. Maliar (2011). Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models. *Quantitative Economics 2*(2), 173–210.

Keefer, M. L., C. C. Caudill, C. A. Peery, and T. C. Bjornn (2006). Route selection in a large river during the homing migration of chinook salmon (oncorhynchus tshawytscha). *Canadian Journal of Fisheries and Aquatic Sciences 63*(8), 1752–1762.

Kelty, M. J. (2005). The role of species mixtures in plantation forestry. *Forest Ecology and Management 233*(2-3), 195–204.

Khoury, C. K., A. D. Bjorkman, H. Dempewolf, J. Ramirez-Villegas, L. Guarino, A. Jarvis, L. H. Rieseberg, and P. C. Struik (2014). Increasing homogeneity in global food supplies and the implications for food security. *Proceedings of the National Academy of Sciences 111*(11), 4001–4006.

Lande, R. (1998). Anthropogenic, ecological and genetic factors in extinction and conservation. *Researches on population ecology 40*(3), 259–269.

Largier, J., B. Magnell, and C. Winant (1993). Subtidal circulation over the northern calofrnia shelf. *Journal of Geophysical Research-Oceans 98*(C10), 18147–18179.

Lee, D. P. and J. Chilton (2007). Hatchery and genetic management plan: American River fall-run Chinook salmon program. Technical report.

Lindley, S. T., C. B. Grimes, M. S. Mohr, W. Peterson, J. Stein, J. T. Anderson, L. Botsford, D. L. Bottom, C. A. Busack, T. K. C. J. Ferguson, J. C. Garza, A. M. Grover, D. G. Hankin, R. G. Kope, P. W. Lawson, A. Low, R. B. MacFarlane, K. Moore, M. Palmer-Zwahlen, F. B. Schwing, J. Smith, C. Tracy, R. Webb, B. K. Wells, and T. H. Williams (2009). What caused the sacramento river fall chinook stock collapse? NOAA-TM-NMFS-SWFSC-447.

Lindley, S. T., R. S. Schick, E. Mora, P. B. Adams, J. J. Anderson, S. Greene, C. Hanson, B. P. May, D. R. McEwan, R. B. MacFarlane, C. Swanson, and J. G. Williams (2007). Framework for assessing viability of threatened and endangered Chinook salmon and steelhead in the Sacramento-San Joaquin basin. *San Francisco Estuary and Watershed Science 5*, 4.

Maliar, L. and S. Maliar (2013). Envelope condition method versus endogenous grid method for solving dynamic programming problems. *Economics Letters 120*(2), 262–266.

McNair, N. N. (2015). An analysis of juvenile chinook salmon outmigration speed and survival in response to habitat features: Sacramento river from knights landing to sacramento, california. University of San Francisco master's thesis, available at http://repository.usfca.edu/thes/143/.

Meyer, M. C. (2008). Inference using shape-restricted regression splines. *The Annals of Applied Statistics 2*(3), 1013–1033.

Meyer, M. C. and X. Liao (2016). Package 'cgam'. R package version 1.5.

O'Farrell, M. R., M. S. Mohr, M. L. Palmer-Zwahlen, and A. Grover (2013). The Sacramento Index (SI). *US Dept. Commerce, NOAA Tech. Memo., NMFS-SWFSC 512*.

Peterson, D., R. Smith, I. Stewart, N. Knowles, C. Soulard, and S. Hager (2005). Snowmelt discharge characteristics sierra nevada, california. *US Geological Survey Scientific Investigations Report*.

Petrosky, C. and H. Schaller (2010). Influence of river conditions during seaward migration and ocean conditions on survival rates of snake river chinook salmon and steelhead. *Ecology of Freshwater Fish 19*(4), 520–536.

Powell, W. B. (2007). *Approximate Dynamic Programming: Solving the curses of dimensionality*. Hoboken, New Jersey: John Wiley & Sons.

Powell, W. B. (2011). *Approximate Dynamic Programming: Solving the curses of dimensionality, second edition*. Hoboken, New Jersey: John Wiley & Sons.

Reisenbichler, R. R. and S. P. Rubin (1999). Genetic changes from artificial propagation of Pacific salmon affect the productivity and viability of supplemented populations. *ICES Journal of Marine Science 56*(4), 459–466.

Rust, J. (1997). Using randomization to break the curse of dimensionality. *Econometrica 65*(3), 487–516.

Schwarzenegger, A. (2008). A proclamation by the governor of the state of California on 10th day of April 2008. State of California, available at https://www.gov.ca.gov/news.php?id=9293.

Simao, H., J. Day, A. George, T. Gifford, and J. Nienow (2009). 2009. *Transportation Science 43*(2), 178–197.

Turelli, M. and N. H. Barton (1994). Genetic and statistical-analyses of strong selection on polygenic traits: What, me normal? *Genetics 138*, 913–941.

United Nations Environment Programme (UNEP) (2005). Millennium ecosystem assessment. *Ecosystems and human wellbeing: a framework for assessment Washington, DC: Island Press*.

United Nations Environment Programme (UNEP) (2007). Global environment outlook: Environment for development (geo4). Nairobi, Kenya: United Nations Environment Programme, available at http://pardee.du.edu/sites/default/files/GEO-4˙Report˙Full˙en.pdf.

Van Roy, B., D. P. Bertsekas, Y. Lee, and J. Tsitsiklis (1997). A neuro-dynamic programming approach to retailer inventory management. *Proceedings of the 36th IEEE conference on decision and control 1-5*, 4052–4057.

Waples, R. S. (1991). Pacific salmon, *Oncorhynchus spp.*, and the definition of "species" under the Endangered Species Act. *Marine Fisheries Review 53*(3), 11–22.

Weitzman, M. L. (1998). The noah's ark problem. *Econometrica*, 1279–1298.

Wells, B. K., J. A. Santora, J. C. Field, R. B. MacFarlane, B. B. Marinovic, and W. J. Sydeman (2012). Population dynamics of chinook salmon oncorhynchus tshawytscha relative to prey availability in the central california coastal region. *Marine Ecology Progress Series 457*, 125–137.

Winship, A. J., M. R. O'Farrell, W. H. Satterthwaite, B. Wells, and M. S. Mohr (2015). Expected future performance of salmon abundance forecast models with varying complexity. *Canadian Journal of Fisheries and Aquatic Sciences 72*(4), 557–569.

Wright, B. and J. Williams (1984). The welfare effects of the introduction of storage. *Quarterly Journal of Economics 99*(1), 169–192.

Zimmermann, F. and C. Jørgensen (2015). Bioeconomic consequences of fishing-induced evolution: a model predicts limited impact on net present value. *Canadian Journal of Fisheries and Aquatic Sciences 72*, 612–624.

Appendix

A Additional model details

A.1 Genetics dynamic equations

The genetic mean value of the trait of interest for juveniles is identical to that of their parents:

$$\mu'_{i,t} = \mu_{i,t} \tag{10}$$

The genetic variance for juveniles is equal to:

$$G'_{i,t} = \frac{\mathcal{V}_{LE}}{2} + \frac{G_{i,t}}{2}, \tag{11}$$

where $G_{i,t}$ is the genetic variance of the returning adult populations and $\mathcal{V}_{LE}$ is the so-called "variance linkage equilibrium". This follows from the random mating of the parents and assumes constant genetic variance at inheritance. Without $\mathcal{V}_{LE}$, selection could eventually drive the genetic variance of each population to zero, which is unrealistic.[25] These measures of central tendency and spread are derived from the dynamics for the full population density distribution over all genotypes of a quantitative genetic trait assuming a normal distribution for each population.

After maturing for a time in the stream, the juvenile salmon out-migrate towards the ocean. This journey involves migration mortality in the stream that is both phenotype-independent (e.g. predation) and phenotype-dependent (i.e. selective) given optimal-timing factors such as stream flow, temperature, etc. Combined migration and selective ocean survivorship (before any ocean harvest) is given by:

$$\kappa_{i,t}\left(D_t|\mu'_{i,t}, G'_{i,t}\right) = \frac{\tilde{\kappa}_i(D_t)\exp\left\{-\dfrac{(E+G'_{i,t})(\varepsilon_t - x_i)^2 + S_i(\varepsilon_t - \mu'_{i,t})^2 + S_o(x_i - \mu'_{i,t})^2}{2\left((E+G'_{i,t})(S_i+S_o) + S_iS_o\right)}\right\}}{\sqrt{\dfrac{(E+G'_{i,t})(S_i+S_o)+S_iS_o}{S_iS_o}}}. \tag{12}$$

[25] We use the infinitesimal model which assumes that a large number of unlinked loci each contribute additively to the overall genotype, such that offspring inherit their genotypes from their mid-parental mean with a variance of $\mathcal{V}_{LE}$ (Turelli and Barton, 1994).

Environmental variance, E, reflects the natural variation that occurs between genotype and phenotype and is fixed across populations.[26] S_j represents selection variance in location j (i.e. in a stream or ocean). Intuition for this parameter comes from noting that $1/S_j$ reflects the strength of the selective mortality at location j. We assume that optimal migration timing with respect to ocean selection is determined by a stochastically drawn upwelling date: $\varepsilon_t \overset{iid}{\sim} \mathcal{N}(\mu_o, \tau_o)$.[27] We assume the migration date selected for in-stream (x_i) is fixed for each stream.

The strength with which in-stream selection acts on the fish population (and thus the strength with streams are naturally kept heterogeneous), $1/S_i$, is dependent on distance trucked for hatchery fish but constant for wild fish:

$$S_i = \begin{cases} S & \text{if } i = 1, 2 \\ S \exp(D_t/\theta) & \text{if } i = 1^*, \end{cases} \tag{13}$$

where θ determines the responsiveness of in-stream selection to trucking. When there is no augmentation, $D_t = 0$, in-stream selection for hatchery fish is the same as for wild fish in stream 1. As trucking distance increases, S_{i*} increases and thus in-stream selection strength for hatchery fish ($1/S_{i*}$) falls.

The mean genetic value (migration date) after both in-stream and ocean selection occur is

$$\mu''_{i,t} = \frac{\mu'_{i,t}(E(S_i + S_o) + S_i S_o) + G'_{i,t}(\varepsilon_t S_i + x_i S_o)}{(E + G'_{i,t})(S_i + S_o) + S_i S_o},$$

and the genotype variance post-selection is

$$G''_{i,t} = \frac{G'_{i,t}(E(S_i + S_o) + S_i S_o)}{(E + G'_{i,t})(S_i + S_o) + S_i S_o}.$$

[26] Environmental variance explains why, for example, individuals with identical genotypes (migration timing genotype) may exhibit different phenotypes (actual migration time) due to differing environmental factors, such as average river flow, flow pulses and shoreline habitat features (McNair, 2015).

[27] We assume that phenotypes are normally distributed around genotypes with environmental variance E, then stabilizing selection acts on those phenotypes for optimal traits μ_x with variance S_x. Specifically, we have two selection events: (1) selection for optimal migration time within each stream (e.g., depending on flow regime), which differs between streams but is constant in time (x_i and S_i), and (2) selection for optimal migration time with respect to ocean upwelling, which is the same across streams and varies in time (ε_t and S_o).

After harvest, fish returning to spawn have a genotype mean given by

$$\begin{aligned}\mu_{1,t+1} &= (1-\sigma)\frac{N'''_{1,t}\mu''_{1,t}}{N_{1,t+1}} + \sigma\frac{N'''_{2,t}\mu''_{2,t}}{N_{1,t+1}} + (1-q(D_t))\frac{N'''_{1^*,t}\mu''_{1^*,t}}{N_{1,t+1}} \\ \mu_{2,t+1} &= \sigma\frac{N'''_{1,t}\mu''_{1,t}}{N_{1,t+1}} + (1-\sigma)\frac{N'''_{2,t}\mu''_{2,t}}{N_{1,t+1}} + q(D_t)\frac{N'''_{1^*,t}\mu''_{1^*,t}}{N_{1,t+1}},\end{aligned}$$

and a genotype variance of

$$\begin{aligned}G_{1,t+1} = (1-\sigma)\frac{N'''_{1,t}(G''_{1,t}+(\mu''_{1,t})^2)}{N_{1,t+1}} + \sigma\frac{N'''_{2,t}(G''_{2,t}+(\mu''_{2,t})^2)}{N_{1,t+1}} \\ + (1-q(D_t))\frac{N'''_{1^*,t}(G''_{1^*,t}+(\mu''_{1^*,t})^2)}{N_{1,t+1}} - (\mu_{1,t+1})^2, \\ G_{2,t+1} = \sigma\frac{N'''_{1,t}(G''_{1,t}+(\mu''_{1,t})^2)}{N_{1,t+1}} + (1-\sigma)\frac{N'''_{2,t}(G''_{2,t}+(\mu''_{2,t})^2)}{N_{1,t+1}} + q(D_t)\frac{N'''_{1^*,t}(G''_{1^*,t}+(\mu''_{1^*,t})^2)}{N_{1,t+1}} \\ - (\mu_{2,t+1})^2.\end{aligned}$$

A.2 Hatchery production and extinction threshold

In our model, the hatchery takes 25% of the returning stock until 600 spawners are selected, such that hatchery production is at its maximum at a stock of 2400 and above. We assume that at 2400, the extinction risk is negligible. In the context of spring- and winter-run Chinook salmon, Lindley et al. (2007) argue that when the stock falls below 250, extinction risk is "high", i.e. the chance of extinction in the next 20 years is over 20%. In our model we use a deterministic threshold for determining when the population goes extinct. We set this threshold at 25, i.e. at 10% of Lindley et al.'s high risk threshold.

A.3 Status quo harvest rule

The Pacific Fishery Management Council's harvest plan for Sacramento River Fall Chinook (SRFC), summarized in Winship et al. (2015, p. 562) is given by

$$F_{SQ}(N_t'') = \begin{cases} \dfrac{0.1N_t''}{\eta_1} & \text{if } N_t'' \in [0, \eta_1) \\ 0.1 & \text{if } N_t'' \in [\eta_1, \eta_2) \\ 0.1 + \dfrac{0.25 - 0.1}{\eta_3 - \eta_2}(N_t'' - \eta_2) & \text{if } N_t'' \in [\eta_2, \eta_3) \\ 0.25 & \text{if } N_t'' \in [\eta_3, \eta_4) \\ \min\left\{1 - \dfrac{\widetilde{N}}{N_t''}, 0.7\right\} & \text{if } N_t'' > \eta_4. \end{cases}$$

The rule specifies constant escapement at $\widetilde{N}$—the stock associated with the maximum sustainable yield—with two exceptions that depend on the aggregate number of all wild and hatchery fish, N_t''. At high stock levels the exploitation rate is capped at $F = 0.7$. At low stock levels, where F would otherwise fall below 0.25, a higher exploitation rate is specified than would result in an escapement of $\widetilde{N}$. $\boldsymbol{\eta} = [\eta_1, \eta_2, \eta_3, \eta_4]$ are the stock levels where the harvest rule changes. From Winship et al. (2015) we have $\boldsymbol{\eta}_{SRFC} = [5e4, 9e4, 11e4, 162.7e4]$. Since we set the scale of our system at one quarter of the size of the aggregate fall run Chinook stock complex, we use $\boldsymbol{\eta} = 0.25\boldsymbol{\eta}_{SRFC}$. We use a smooth, best-fit approximation to the rule given by:

$$F(N_t'') = \min\{7.7025\text{e}{-6} \cdot N_t'', 0.702\}. \tag{14}$$

A.4 Profit function parameters

There exists relatively good information on California chinook harvest revenues but relatively poor information on harvest cost. Information on stock levels, harvest levels (numbers and weights), and prices is available from the "Stock Assessment and Fishery Evaluation Document for the Pacific Coast Salmon Fishery Management Plan" (PFMC, 2014). We set p, the price per fish, at \$76.9 given by the product of the 2010-2014 average pounds of dressed weight per fish (13.4 lbs) and price per pound (\$5.74 per lb). When there are no ready substitutes for fish consumed from a particular stock, we would expect the unit price to generally increase as the quantity supplied in a given year falls. However, for California chinook prices from 2006-2014 have been relatively stable (\$5.44-\$6.33 per lb) despite wide swings in levels of commercial harvest (0.23M - 3.79M lbs). Because there is no discernible trend in the price as quantity varies, we assume that p is constant over the level of harvest. In a model of the

California chinook commercial fishery, Garnache (2015) parameterizes cost by calibrating to a total harvest cost estimate for the fishery from 2006. Because 2006 involved relatively low harvest and high escapement, this approach likely leads to a high estimate of costs. Indeed, such a parameterization implies that any exploitation rate above approximately 36% results in negative profits. Observed exploitation rates in this system routinely exceed this level. We parameterize the cost function such that, when the harvest control rule is considered, profits do not decline as the stock level increases ($c = 39.8$).

A.5 FDP solution method details

The number of periods per simulation, T, is selected to balance the tradeoff between ensuring the implications of starting at a given state are "felt" (e.g. the potential for extinction) and avoiding excessive representation of the steady state region in which simulation chains congregate given sufficient time. The number of simulations in a block between each regression step, z, is selected to balance another tradeoff, i.e. between better representation across the state space and slower incorporation of new information and movement towards convergence (as z increases).

The step size, δ_t, specifies the relative weighting of new versus existing information. A larger weight on new information is useful at the outset since the initial guess may be poor. However, as information accumulates in the value function estimate a lower weight is desirable to temper the influence of new stochastic realizations of value. A number of alternative step size functions have been explored, including constant and decreasing weights (see Powell, 2011). We make a minor contribution to FDP via this component by specification of the following hybrid approach. We use a step size that is relatively high and constant while the value function estimate is moving consistently towards higher or lower values. Once the value function estimate stabilizes in magnitude[28] we switch to a step size function that decreases exponentially until reaching a lower bound:

$$\delta_t = \max\{\alpha \exp(-\gamma(z - z_s), 0.05\}, \tag{15}$$

where α is the initial and maximum weight given to new information, γ is the rate at which the weight decays as the number of regression steps increase, z is the regression counter, and z_s is the counter value when the switch to the declining function occurs. This hybrid approach facilitates an initially high weight on new information for aggressive updating, with subsequent low weight to facilitate convergence by tempering the effect of stochastic

[28]Specifically, we calculate the following. Define $E\bar{V}^z = E[\bar{V}^z(\hat{X})]$ as the mean value function given a probability mass function over states as observed over approximately the last 50 regression blocks . For each regression block z we calculate the absolute relative deviation in the mean value, $\text{abs}\{\left(E\bar{V}^z - E\bar{V}^{z-1}\right)/E\bar{V}^z\}$. If the value function is no longer consistently iterating towards higher or lower values, we would expect this measure to be small. We initiate the switch to a decreasing step size when the average of this statistic over the last 5 regressions is less than $\phi = 0.001$. We consider the average over multiple periods to avoid idiosyncratically triggering the switch when two regression models in a sequence are similar by chance and not due to stabilization.

realizations of value.

In our model for the value function, each state variable enters additively in nonparametric form. We also include several interaction terms but do so parametrically for two reasons. First, when variables are interacted (multiplied) this can induce large gaps between data points which can cause nonparametric estimation to falter. Second, parametric specification provides coefficient estimates that help in unpacking complex interactions. We might expect the value of biomass in one stream to depend on the other and therefor include an interaction between N_1 and N_2. It is also likely that the value of genetic variance will depend on the genetic mean and therefor we include interactions between μ_i and G_i for both streams. Finally we also include $\mu_i^2 G_i$ terms for both streams since the value of the genetic mean is likely to have an interior peak (i.e. extreme high and low values for the genetic mean are likely to be disadvantageous). The parametric contribution to the value function is specified by: $\rho_0 + \rho_1 N_1 N_2 + \rho_2 \mu_1 G_1 + \rho_3 \mu_2 G_2 + \rho_4 \mu_1^2 G_1 + \rho_5 \mu_2^2 G_2$.

A.6 Summary of model parameters

In Table 2 we summarize the model parameters.

	Description	Value	Units
ζ	fraction of stream 1 returning adults that the hatchery takes	0.25	
H_{max}	maximum number of adults the hatchery can take	$600^{\dagger}$	fish
R_i, $i = 1, 2, 1^*$	fecundity	200	smolt per spawner
K_i, $i = 1, 2$	carrying capacity of stream i	3×10^5	smolt
$\mathcal{V}_{LE}$	variance linkage equilibrium	5	
κ_{wild}	wild in-stream survival, not selection based	0.7	
E	environmental variance	10	
$\underline{N}$	minimum viable population	25	fish
S_i, $i = 1, 2$	inverse in-stream selection strength	20	
S_o	inverse ocean selection strength	20	
(μ_1, μ_2)	in-stream selection for migration in streams 1 and 2	$(43, 57)^{\dagger}$	Julian day
(μ_o, τ_o)	mean and variance of ideal migration day, with respect to ocean upwelling	$(50, 25^{\dagger})$	Julian day
M	instantaneous rate of natural mortality	0.2	
θ	responsiveness parameter for in-stream selection to augmentation	0.5	
σ	natural straying rate	0.05	
q_{max}	straying rate if trucked the maximum distance	0.5	
p	price per fish	76.9	\$
c	harvest cost parameter	39.8	\$
c_{truck}	cost of trucking maximum distance	200	\$
β	discount factor	1/1.03	
α	initial step size	0.85	
γ	step size decay rate over number of regression steps	5×10^{-5}	
ϕ	threshold for switching to decreasing step size (absolute relative deviation in the mean value)	1×10^{-3}	
m	number of regression updating steps over which the convergence metric is averaged	10	
ω	tolerance for convergence	2.5×10^{-3}	

Table 2: Summary of parameters. Biological model parameter values used are the same as in Dedrick and Baskett (2018), except where noted with a †.

B Additional results

B.1 Additional value function results

In Figure 9 we present the value function versus the genetic variance (G_i) of stock 1 and 2 (rows), respectively, for a range of mean migration dates (columns). Two key results emerge: the effect of genetic variance on the value function is small and is least beneficial (even detrimental) to value when the genetic mean is near its ideal. We exclude the upper range of G_i in the figure since these levels are exceedingly rare (discussed further below). The effect on the value function of increasing G_i for either stream from its lower to upper bound in Figure 9 is no more than 1%. When μ_i is near its ideal for stream 1 ($\mu_1 = 49.5$, row 1, column 2) the benefit to increasing G_1 is at its lowest. For stream 2, when μ_2 is near its ideal ($\mu_2 = 52.7$, row 2, column 3) the value actually falls as G_2 increases. Overall, we find

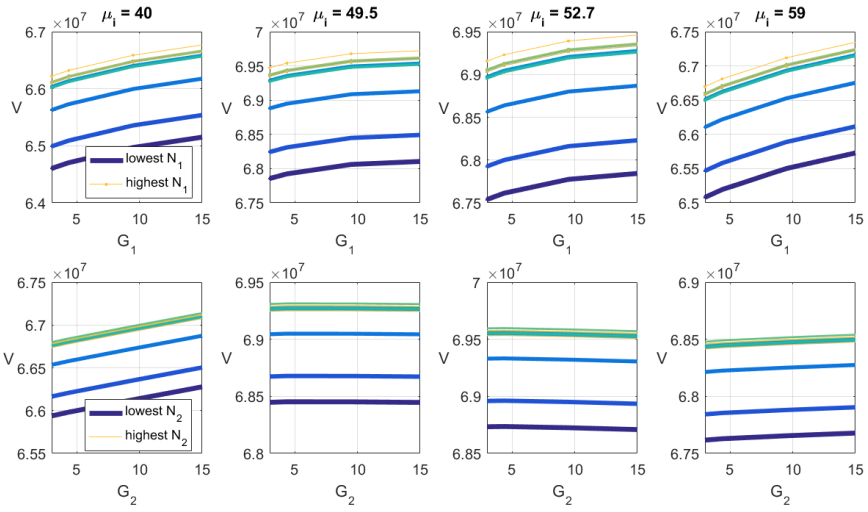

Figure 9: The value function over the trait variance state (G_i) in stream 1 (top row) and stream 2 (bottom row) for various levels of the mean trait state (μ_i, columns). Multiple curves within each panel depict the value at various stock levels (N_i), from lowest (thickest line) to highest (thinnest line). All states that do not vary in a panel are set to their modal levels.

that when the genetic mean is best adapted to the conditions faced by a population (stream and ocean) variance is of little (and possibly negative) value. This latter result is intuitive but the weak value of variance in general is surprising since it is conventionally believed to be quite valuable, at least when when the population is not well adapted.

The value function model is nonparametric except for a set of additive, parametric interaction terms. For additional stability in regression we model a standardized value function where all variables have been rescaled to have a mean of zero and variance of one. Coefficients and standard errors are presented in Table 3.

constant	0.00141
	(0.00068)
$N_1 N_2$	-0.047
	(0.0012)
$\mu_1 G_1$	-20.5
	(0.16)
$\mu_2 G_2$	-12.92
	(0.15)
$\mu_1^2 G_1$	10.98
	(0.083)
$\mu_2^2 G_2$	6.37
	(0.081)

Table 3: Regression coefficients and standard errors (in parenthesis) for parametric interaction terms in value function model. All coefficients are significantly different from zero ($p < 0.01$).

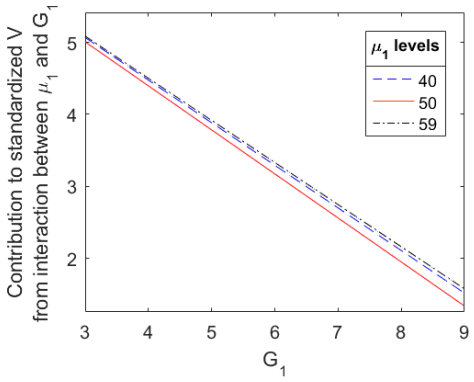

Figure 10: Contribution to standardized V (mean zero, unit variance) from the interaction between μ_1 and G_1 as a function of G_1.

Figure 10 shows the contribution to the standardized value function from the interaction between μ_1 and G_1 as a function of G_1. The figure does not include the direct effect on value from μ_1 and G_1. This figure illustrates that, in general, as G_i increases the value falls faster when μ_i is near its ideal.

B.2 Additional simulation results

Figure 11 shows the temporal mean of levels for each state variable and for profit. Figure 12 shows the relative state variable frequencies from Monte Carlo simulations given no augmentation and the optimal policy. In the final column we see a negative correlation between genetic variance in the two streams because ideal ocean arrival conditions that shift towards timing that favors one stream reduces drivers for spread in that stream while increasing spread in the alternative stream. This dynamic is disrupted under optimal augmentation. Table 4 shows the summary statistics for state variables from the Monte Carlo simulations under the optimal policy after excluding burn-in periods.

	N_1	μ_1	G_1	N_2	μ_2	G_2
Mean	18273	47.8	3.85	14820	48.8	4.48
Mode	24314	47.7	3.75	14744	48.7	4.56
Std dev.	5801	1.1	0.17	3369	1.3	0.24

Table 4: State variable statistics from simulations under the optimal policy after excluding burn-in periods.

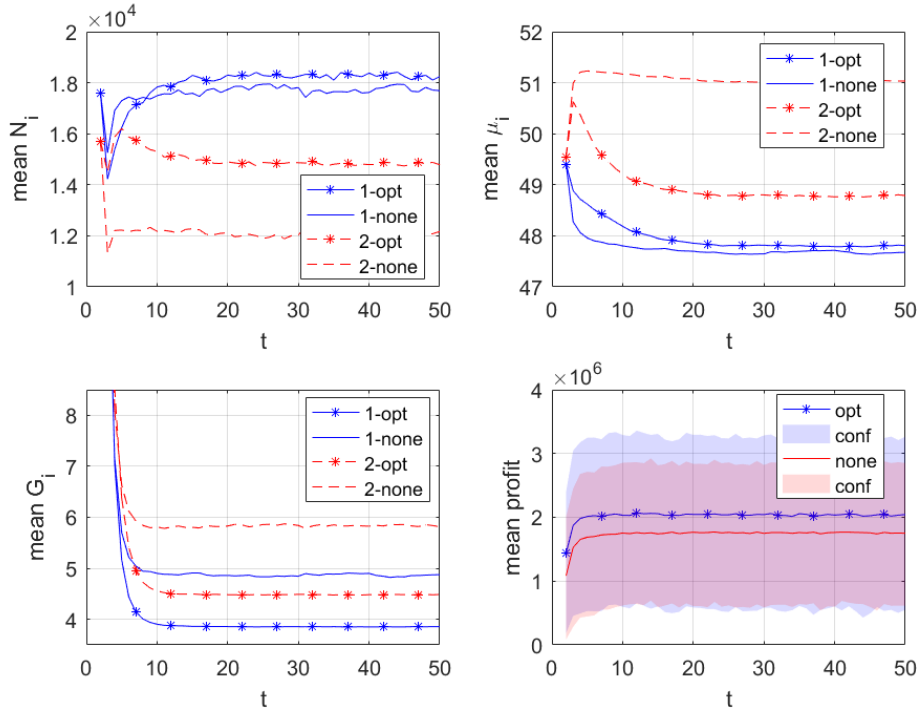

Figure 11: Mean levels for state variables (for streams 1 and 2) and profit under optimal ('opt') and no ('none') augmentation over 3,000 simulation runs across 50 periods. In the final panel, 90% confidence intervals ('conf') are shown around mean profit.

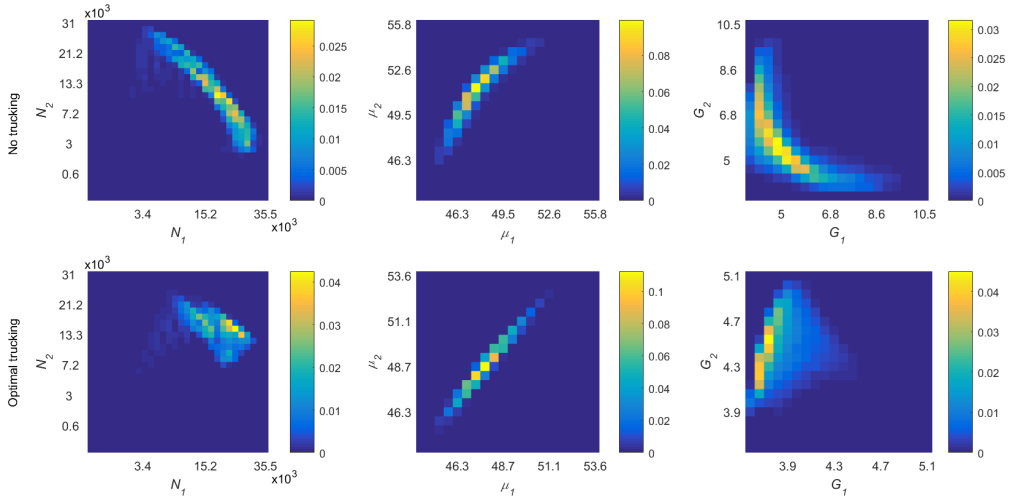

Figure 12: Relative state variable frequencies in simulations under no augmentation (top row) and the optimal policy (bottom row), for stream 1 (horizontal axis) and stream 2 (vertical axis), for each state type, (N, μ, G). Outcomes reflect 3,000 simulations, starting from randomly selected loci in the state space, of 50 periods each, with the initial 30 periods excluded as burn-in periods. Continuous state variable values were binned to the nearest node for the plot.