Title

Using Ecological Models to Inform Decisions in Dynamic and Uncertain Environments

Author

Buckner, Jack Hoyt

Publication Date

2023

Peer reviewed|Thesis/dissertation

Using Ecological Models to Inform Decisions in Dynamic and Uncertain Environments

By

JOHN H BUCKNER
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Ecology

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Marissa L. Baskett, co-Chair

Michael R. Springborn, co-Chair

James N. Sanchirico

Committee in Charge

2023

Contents

Abstract

I present results from three studies addressing apparently disparate challenges: conserving biodiversity, maintaining sustainable fisheries, and controlling infectious disease. All, however, require fateful choices to be made urgently in the context of substantial uncertainty and environmental change. Uncertainty and environmental change can have complex interactive effects on the decision-making process. Uncertainty adds an element of risk that can be amplified over time by changing conditions and dynamics feedback mechanisms. Furthermore, the combination of uncertainty and time introduces the possibility of learning; policymakers may choose to delay fateful choices until more information is available, invest in improved knowledge, and take experimental actions. For my dissertation, I developed mathematical models to understand how these interactive effects shape optimal decision-making in three applied case studies: prioritizing scarce COVID-19 vaccines, mitigating unintended genetic impacts of fish hatcheries, and designing fisheries monitoring programs. Across the case studies, I found that the time scales of underlying biological processes were critical for designing effective management strategies. Furthermore, when decision-makers respond adaptively to changes in the underlying ecosystem state, they can act as a stabilizing feedback mechanism, reducing the risks introduced by stochastic environmental variability and scientific uncertainty.

Acknowledgments

I want to thank my advisors, Michael Springborn and Marissa Baskett, for their guidance, help, and encouragement. As well as the Baskett and NatuRE Policy lab groups at UC Davis for providing fun communities to engage with.

I should also acknowledge the funding source for my dissertation research. I was very fortunate to be supported by the National Science Foundation (NSF) Graduate Research Fellowship, which gave my a huge amount of autonomy over my research program and opportunities to develop new skills. I also participated in the NSF-funded Sustainable Oceans training program, which taught me a lot about approaching policy-focused research and helped me establish important professional connections. Finally, the British Columbia Ministry of Forests funded the work in my second chapter. I am grateful for the opportunity to work on applied conservation issues with scientists from the provincial government and the funding that supported this work.

Finally, none of this work would have been possible without the love and support of my family, Jane, Fred, and Emily Buckner, and my fiance, Rebecca Fairchild.

CHAPTER 1

Introduction

This dissertation presents three studies addressing apparently disparate challenges: conserving biodiversity, maintaining sustainable fisheries, and controlling infectious disease. All, however, require fateful choices to be made urgently in the context of substantial uncertainty and environmental change. The changes these systems exhibit are dynamic: multiple components of the system simultaneously influence one another in cycles of mutual causation. These causal loops are called feedback mechanisms, and they can amplify or dampen the impacts of our decisions and produce to complex emergent behaviors in deceptively simple systems (**1**, **2**). In the context of decision-making, these feedback mechanisms cause choices in the present to have ramifying consequences that influence the options available in the future (**3**).

Uncertainty is a more intuitive concept than dynamic change; it boils down to having incomplete information about what the consequences of our actions will be, although it can arise in a variety of different forms (**4**, **5**). Despite this simple definition, when combined with dynamic environmental change, it can have complex and counter-intuitive implications for decision-making. Feedback mechanisms can amplify uncertainty; for example, small variations in biological processes, like the reproductive number of a disease, can determine the difference between an epidemic petering out or growing into a pandemic (**6**). Furthermore, the combination of uncertainty and time introduces the possibility of learning, adding a new layer to the decision-making process. Policymakers may take active measures to reduce uncertainty through research and monitoring (**7**), delay a fateful choice until more information is available (**8**), or take a chance on a new idea to learn how well it works (**9**).

The goal of my research, in general and in this dissertation in particular, is to understand the interactive effects of uncertainty and dynamic biological processes on the design of environmental policy. I approach this general question by developing mathematical models to analyze specific applied decision problems. This emphasis on specific case studies is necessary because the effects of dynamic feedback mechanisms and uncertainty on decision-making are context-dependent: the optimal approach to management depends on specific features of the study system and the goals managers want to achieve. As a consequence, results from one case study are unlikely to hold general. However, working on these specific case studies can help motivate the development of novel mathematical tools and conceptual frameworks that can be more widely applicable. Given these constraints, my research process is a dialectic between working on specific applied problems

using conceptual frameworks to generate hypotheses and refining these frameworks based on the findings from the case studies.

Using this approach, I present papers from three case studies largely chosen for their applied relevance, although each leverages mathematical and computational tools and incorporates uncertainty and dynamic biological processes. Chapter 2 presents results from a non-linear optimization model used to identify strategies for allocating scarce COVID-19 vaccines when they became available in the spring of 2021. The strategies are influenced by feedback between infection, vaccination, the number of susceptible individuals in the population, and uncertainty about the rates of virus transmission between demographic groups. Chapter 3 uses a population genetic model to identify strategies to mitigate unintended genetic impacts of a hatchery program used to maintain a threatened population of White Sturgeon in British Columbia. Effective mitigation strategies are determined by feedback between the population's demography and the distributions of genotypes within the population; the total magnitude of genetic effects is also difficult to measure in long-lived species, creating uncertainty about the need for mitigation measures in this system. In Chapter 2, I develop a bioeconomic model to characterize the optimal strategy for investing in monitoring to reduce uncertainty about the abundance of harvested fish populations. This model uses tools from decision theory to link the information from a monitoring program to the economic value of the decisions it informs.

In all three projects, one can consider the relevant decision in terms of coupled human-natural systems. The basic premise of coupled human natural system analyses is understanding the mutual causal links between human actives and their environment and the outcomes generated by the feedback between these components of the system (**10**). Coupled human-natural systems analysis can be applied to a range of environmental and human systems, but to analyze decision-making, I focus on the links between a decision maker and how their choices affect the environment and how the decision-maker responds to these changes. A diagram of the relevant causal links is given in Figure 1.1. The framework has three primary components 1) the environment, 2) human activities that impact the environment, and 3) a manager regulates the level of human impacts. The manager makes decisions based on observations state of the decision of the environment and their understanding of the dynamics of the environmental system. In turn, these choices (imperfectly) influence the level of human impacts on the environment, creating a feedback mechanism.

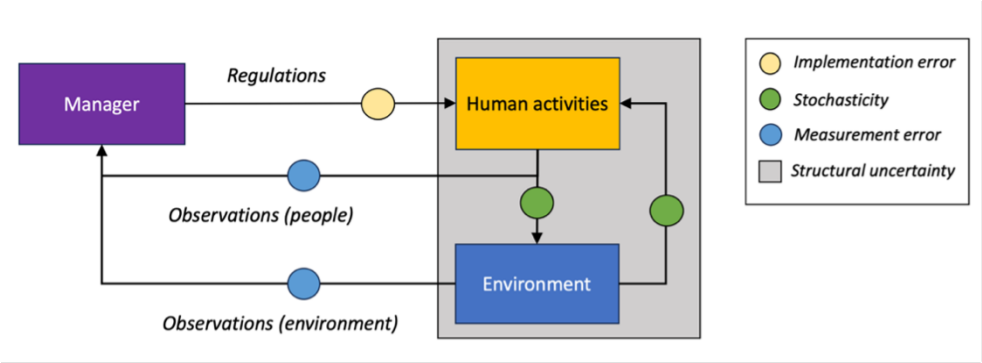

FIGURE 1.1. Decision-making under uncertainty in the context of coupled human-natural systems. Boxes show the interacting components of the system with black arrows indicating causal links between them and colored dots representing the effect of uncertainty on these links. I assume the model includes a description of some human activities (yellow) that affect the state of the environment (blue), such as resource extraction, restoration, etc. The state of the environment can in turn, influence the magnitude of human impacts. The manager adds in direct feedback between the environment and human activities. Information about the state of the environment affects their choice, which in turn influences the magnitude of human impacts on the system.

Thinking about the decision maker as a component of a larger system highlights a key insight: echoing Hastings (2016), the time scale of the environmental process is critical for designing effective management strategies (**11**). Many modeling studies solely focused on biological processes will ignore the time scale of dynamic processes by studying equilibrium behaviors. Because the manager and the environmental processes are embedded in the same dynamical system, however, the relationship between the time scales of the biological and management processes can determine the effectiveness of management strategies. For example, in Chapter 2, the model describes how limited supply reduces the rate COVID-19 vaccines can be distributed to the population. We found that the rate of transmission compared to the rate of vaccine distribution determined the optimal prioritization strategy. When vaccines could be administered quickly relative to disease transmission strategies that targeted demographic groups who were most likely to become infected and spread infections to others, minimized total deaths from infection. In contrast, when supply was limited, targeting demographic groups with the highest mortality rate minimized total deaths. In Chapter 4, we find that the long-lived life history of white sturgeon causes the genetic impacts of hatchery production to accumulate slowly over time. The slow time scale of the biological dynamics facilitated adaptive management strategies that require time to learn and improve. This reduced

the risks of operating the hatchery program in the short run because it provided a long window of time to reduce uncertainty and test strategies to minimize the genetic impacts of the hatchery. Finally, in Chapter 4, we found that under the optimal monitoring strategy, the frequency of monitoring was largely determined by the time scale of the population dynamics. More productive populations with higher turnover rates were monitored more frequently under the optimal policy. This tightened the feedback between changes in the population's abundance and the manager's decisions, increasing the stability of the system.

In addition to demonstrating the importance of time scales in environmental management, the coupled human-natural system framework demonstrates that decision-makers function as a feedback mechanism. By using information about the environment to adjust the level of human impacts they can increase the resilience [1] of ecological systems (**12**). Although resilience is not always a desirable characteristic, it can improve outcomes in resource management problems that involve maintaining the environment in a state where it produces valuable ecosystem services. Fisheries management (the topic of Chapter 4) is an instructive example; in the single species context, fish stocks are most productive at intermediate levels of biomass, where it has the highest growth rate (the maximum sustainable yield). In this context, management strategies that maintain the stock in this productive state (i.e. promote resilience) produce the largest amount of harvest. In fact, the strategy that maximizes the harvest from a fishery also maximizes the resilience [2] of the stock around the optimal equilibrium biomass [3]. In Chapter 4, I show that the value of monitoring for fisheries management derives its value from its role in creating a stabilizing feedback mechanism between the stock's abundance and harvest.

[1] Resilience, annoyingly, has a few definitions. In this case, I mean the general propensity of a system to return to a steady state after a disturbance.

[2] Here I mean engineering resilience: the rate of return to the steady state.

[3] Yes, this is true! The strategy that maximizes discounted harvest is a "bang bang" solution (Reed 1979, Clark 1993). The optimal strategy fishes the stock down to an optimal equilibrium level that depends on the discount rate and the productivity of the stock. When the abundance of the stock is above this value, the maximum harvest rate is applied, causing the stock to return as quickly as possible to the equilibrium abundance. When the stock is below the equilibrium point, harvest is set at the minimum value to maximize the rate it returns to the equilibrium.

Box 1: Adaptive policies outperform static ones under structural uncertainty and resource stochasticity

To illustrate this point, I used a simple model of a fishery that tracks the biomass of the stock B and harvest H over time. The stock has density-dependent growth with growth rate r and carrying capacity K; stochasticity was included in the model with multiplicative growth shocks described with a Brownian motion term with variance σ

$$dB = rB\left(1 - \frac{B}{K}\right)dt + \sigma B dW \tag{1.1}$$

Following Tilman et al. (2023), in the absence of management, the harvest level is determined by the abundance of the stock to reflect the fact that anglers reduce fishing effort when the stock is scarce

$$H = \frac{cB^2}{h^2 + B^2} \tag{1.2}$$

where c and h are constants that determine the shape of the relationship between biomass and harvest (**13**). I chose the values of these constants so that the stock would become overfished in the absence of management.

The manager chooses a harvest limit $\overline{H}$, which I assume acts as an upper bound on the harvest rate $H < \overline{H}$. I compare two strategies for setting the harvest limit: 1) a constant strategy where $\overline{H}$ is constant over time and 2) a responsive strategy where $\overline{H}$ is set to a fraction f of the stock's biomass $\overline{H}(B) = fB$. For these two policies, I chose the values of $\overline{H}$ and f to achieve 95% of the maximum sustainable yield (MSY) of the stock.

I analyzed the performance of these two strategies with simulations. To illustrate the effects of structural uncertainty, I simulated the model with the correct value of the harvest limit $\overline{H}$ and with a 10% error above and below the true value. The responsive policy stayed near the maximum biomass that produces the maximum sustainable yield ($B_{MSY} = 1$) in all three cases (figure 1.2.A blue). However, the stock became overfished under the constant harvest limit when the productivity of the stock was lower than expected (figure 1.2.A orange).

To model stochastic population dynamics, I set the variance term of the model to $\sigma = 0.05$. The results were very similar to the structural uncertainty case, with the responsive policy staying near B_{MSY} and the constant harvest policy resulting in overfishing in some cases (figure 1.2.B).

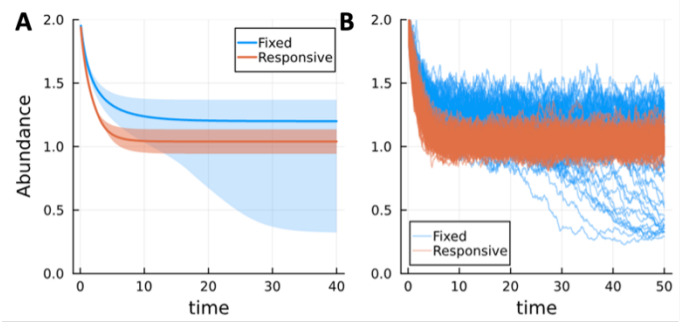

Figure 1.2. Comparing the performance of policies that incorporate feedback with the environment (orange) to fixed policies (blue), under conditions of structural uncertainty (A) and resource stochasticity (B). parameter values: $r = 1.0$, $K = 2.0$, $c = 4.5$, and $h = 7.0$

Using management as a stabilizing feedback mechanism can improve outcomes in dynamic environment systems under multiple forms of uncertainty. Box 1 illustrates this point with a simulation from a simple model of fishery. A more thorough analysis of similar policies has also demonstrated

that incorporating stabilizing feedback improves outcomes in the context of directional environmental changes caused by global warming (**14**). By adjusting management in response to observed changes in the state of the system, responsive management strategies automatically incorporate information about the system's biology over time, reducing their sensitivity to initial assumptions and uncertainties. In effect, incorporating feedback mechanisms is a simple form of passive adaptive management (**5**, **15**). Analyzing policies as feedback mechanisms might be a useful heuristic for designing adaptive management strategies in systems that are too complex for formal optimization tools (**16**).

Designing environmental policy has always been challenged by the prevalence of uncertain and dynamic environmental change, and these complicating factors are playing an increasing role as the climate change. These problems can be addressed with both quantitative models as well as conceptual models, and theoretical frameworks. My dissertation research makes specific and general contributions in both dimensions. In the future, I am interested in building on these insights with further research into the role of management as a feedback mechanism in coupled human-natural systems and developing models to study the decision-making under uncertainty in environments undergoing directional or non-stationary change that are more reflective of the impacts of global warming on natural resource systems.

Bibliography

1. E. N. Lorenz, *Journal of the Atmospheric Sciences* **20**, 130 (1963).
2. A. Hastings, K. Higgins, *Science* **263**, 1133 (1994).
3. C. Clark, *Mathematical Bioeconomics: The Optimal Management of Renewable Resources.*, vol. 49 (Wiley, 1993).
4. H. M. Regan, M. Colyvan, M. A. Burgman, *Ecological Applications* **12**, 618 (2002).
5. J. LaRiviere, D. Kling, J. N. Sanchirico, C. Sims, M. Springborn, *Review of Environmental Economics and Policy* **12**, 92 (2018).
6. F. Brauer, *Infectious Disease Modelling* **2**, 113 (2017).
7. P. L. Fackler, R. G. Haight, *Resource and Energy Economics* **37**, 226 (2014).
8. K. J. Arrow, A. C. Fisher, *Classic papers in natural resource economics* (Springer, 1974), pp. 76–84.
9. C. J. Walters, R. Hilborn, *Journal of the Fisheries Research Board of Canada* **33**, 145 (1976).
10. Y. A. Shin, K. Lacasse, L. J. Gross, B. Beckage, *Ecology and Society* **27**, art4 (2022).
11. A. Hastings, *Proceedings of the National Academy of Sciences* **113**, 14568 (2016).
12. S. A. Levin, J. Lubchenco, *BioScience* **58**, 27 (2008).
13. A. R. Tilman, E. H. Krueger, L. C. McManus, J. R. Watson (2023). Publisher: arXiv Version Number: 1.
14. J. P. Kritzer, C. Costello, T. Mangin, S. L. Smith, *ICES Journal of Marine Science* **76**, 1424 (2019).
15. M. Springborn, J. N. Sanchirico, *Journal of Environmental Economics and Management* **66**, 609 (2013).
16. P. Fackler, K. Pacifici, *Journal of Environmental Management* **133**, 27 (2014).

CHAPTER 2

Dynamic Prioritization of COVID-19 Vaccines When Social Distancing is Limited for Essential Workers

2.1. Author's note

My co-authors, Michael Springborn and Gerardo Chowell, and I wrote this paper in the winter of 2020. It was originally published in the spring of 2021 in the Proceedings of the National Academy of Sciences. The original paper is reproduced here, some of the discussion will reflect conditions at the time of writing and publication.

2.2. Abstract

COVID-19 vaccines have been authorized in multiple countries and more are under rapid development. Careful design of a vaccine prioritization strategy across socio-demographic groups is a crucial public policy challenge given that (1) vaccine supply will be constrained for the first several months of the vaccination campaign, (2) there are stark differences in transmission and severity of impacts from SARS-CoV-2 across groups, and (3) SARS-CoV-2 differs markedly from previous pandemic viruses. We assess the optimal allocation of a limited vaccine supply in the U.S. across groups differentiated by age and also essential worker status, which constrains opportunities for social distancing. We model transmission dynamics using a compartmental model parameterized to capture current understanding of the epidemiological characteristics of COVID-19, including key sources of group heterogeneity (susceptibility, severity, and contact rates). We investigate three alternative policy objectives (minimizing infections, years of life lost, or deaths) and model a dynamic strategy that evolves with the population epidemiological status. We find that this temporal flexibility contributes substantially to public health goals. Older essential workers are typically targeted first. However, depending on the objective, younger essential workers are prioritized to control spread or seniors to directly control mortality. When the objective is minimizing deaths, relative to an untargeted approach, prioritization averts deaths on a range between 20,000 (when non-pharmaceutical interventions are strong) and 300,000 (when these interventions are weak). We illustrate how optimal prioritization is sensitive to several factors, most notably vaccine effectiveness and supply, rate of transmission, and the magnitude of initial infections.

2.3. Introduction

As the novel coronavirus (SARS-CoV-2) continues to inflict substantial morbidity and mortality around the world despite intervention efforts, public health experts see vaccines as essential to

dramatically reduce the mortality burden and possibly halt local transmission (**1**). Novel coronavirus disease 2019 (COVID-19) has resulted in over 2.3 million confirmed deaths globally (**2**) as of early February 2021. Fortunately, multiple promising vaccines are under rapid development, with the final weeks of 2020 seeing the first authorization and shipping of doses (**3**). However, vaccine availability will be highly constrained for at least several months (**4**). This scarcity, combined with stark differences in the spread and impact of SARS-CoV-2 across demographic groups, means that vaccine prioritization poses a key public health challenge. National and international public health organizations have mobilized to assemble guidance, including the World Health Organization (WHO), the National Academy of Medicine (NAM), and the Advisory Committee on Immunization Practices (ACIP) of the US Centers for Disease Control and Prevention (CDC) (**5**).

An effective public health policy for pandemic vaccine allocation requires an understanding of how risk of infection and severe disease varies across socio-demographic groups and how a given vaccine policy will impact the continued spread of infections within the population. Accounting for these two processes is critical when the population with the greatest risk of infection differ from those with the greatest risk of severe disease, as is the case for COVID-19, because an effective policy will need to balance direct protection of the most vulnerable against limiting secondary infections and rapidly achieving herd immunity (**6**). These key components can be integrated into a mathematical and statistical modeling framework of the transmission dynamics of the novel pathogen. Such an analytic framework can then be utilized to investigate the optimal vaccine allocation strategies to achieve a defined public health objective while taking into account the value of vaccines for mitigating health outcomes at the individual and population level.

Previous experience with vaccine development mid-pandemic offers limited insights for SARS-CoV-2 prioritization. SARS and Zika vaccine development was incomplete when those outbreaks ended (**7**). In 2009, as the novel A/H1N1 influenza virus continued to spread across the U.S., researchers investigated optimal vaccination strategies using an age-structured dynamical model. They found that school-aged children and their parents should be prioritized, a strategy that would indirectly protect individuals at higher risk of severe health outcomes (**8**). Sharp differences in the epidemiology of human influenza and COVID-19 indicate that vaccination strategies against the ongoing pandemic should not simply mirror vaccination policies against influenza. For example, COVID-19 is associated with lower susceptibility to infection among children and adolescents (**9**,**10**) and has a

substantially higher infection fatality rate overall that also increases markedly with age (**11**). Toner et al. (**5**, p. 24) provide a detailed overview of the 2018 pandemic influenza vaccination plan and conclude that, "the priority scheme envisioned...does not comport with the realities of the COVID-19 pandemic and new guidance is needed." Fitzpatrick and Galvani (**12**) concur, detailing how the unique "epidemiological, clinical, behavioral, and vaccine-related relationships" of SARS-CoV-2 motivate the need for "pathogen-specific transmission modeling".

We develop and apply a mathematical model to assess the optimal allocation of limited COVID-19 vaccine supply in the U.S. across socio-demographic groups differentiated by age and essential worker status (see Methods). The transmission dynamics are modeled using a compartmental model tracking eight demographic groups through the nine disease states as shown in Fig. 2.1. The parameters are calibrated to capture our current understanding of the epidemiology of COVID-19, and our analysis is designed to capture two key features of COVID-19 prioritization: essential workers and the gradual availability of vaccines over time. A large number of workers are constrained in their ability to work from home (essential workers) exposing them to higher level of risk of infection, and increasing the chance they transmit the disease if infected. Policies that account for the greater risk essential workers are exposed to may be more just and highlight a group of individuals "who have been overlooked in previous allocation schemes" (**5**). Furthermore, these policies may be more effective at mitigating morbidity and mortality as they can account for a key factor driving transmission of the disease.

To account for the gradual rollout of vaccines, we employ stochastic non-linear programming techniques to solve for vaccine prioritization policies that distribute vaccines to susceptible individuals and change on a monthly time step responding to changes in the epidemiological status of the population (shares of the population in different disease states). These dynamic policies account for a key feature of the policy-making process since the supply of vaccines is likely to be constrained with available doses administered as they become available over a period of several months.

The transmission of COVID-19 is a complex process contingent on the characteristics of the disease and ever changing social behavior. Furthermore many of the key dynamics can change depending on the spatial scale considered, with differences in the transmission process within and between communities. We seek to summarize the features of the complex and evolving processes that are most relevant to the spread of the disease within and between socio-demographic groups. To do so

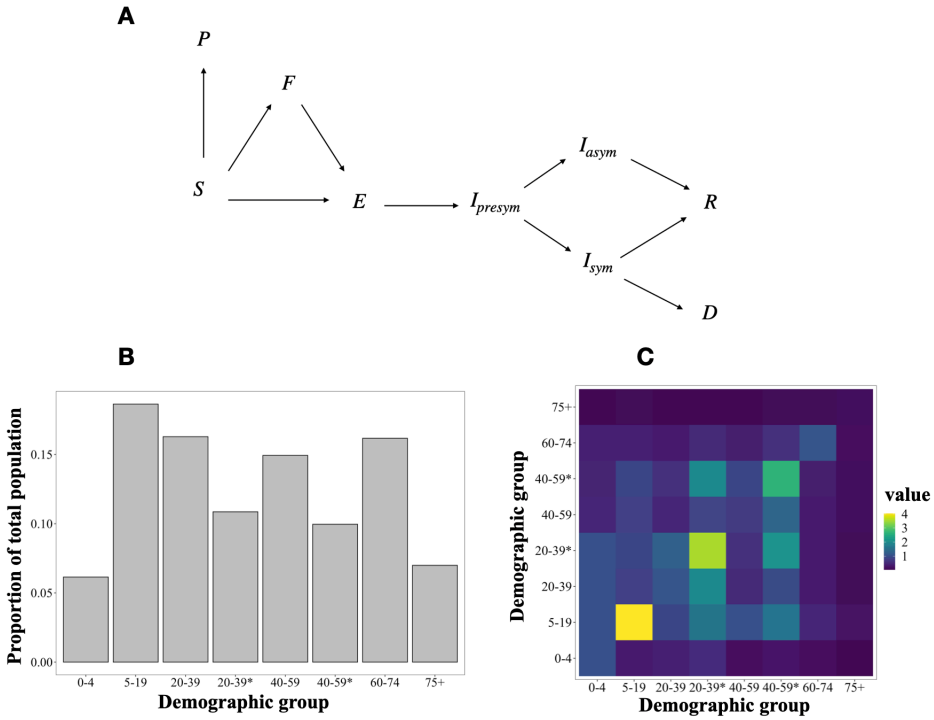

FIGURE 2.1. Schematic of the modeled movement of individuals between epidemiological states defined in Methods (A), the portion of individuals from the U.S. population in each demographic group determined by essential worker status (*) and age (B), and the contact rates between demographic groups, given by average daily number of contacts a group on the horizontal axis makes with a group on the vertical axis (C).

we model COVID-19 transmission with the social contact hypothesis (**13**) and describe the contact patterns between demographic groups using contact matrices estimated for the United States from Prem et al. (**14**) scaled by the location where the contacts were made (home, school, work and other) to reflect the impacts of social distancing. Although these assumptions present a stylized version of contacts during the pandemic, they allow us to capture many key features of social contacts, such as the concentration of contacts within age groups, parent-child relationships and receiver-caregiver relationships (**15**).

Existing published studies of COVID-19 vaccination prioritization analyses include Matrajt et al. (**16**) and Bubar et al. (**17**).[1] Both consider the optimal allocation of vaccines across five or more age groups within a country. Their approaches feature rich exploration of policy sensitivity to vaccine effectiveness and availability. Matrajt et al. is particularly detailed in this respect,

[1]We also note (**18**) use simulation without optimization to explore implications of vaccines with various levels of direct and indirect protection.

while Bubar et al. uniquely consider differences in demographics and contact rates across multiple countries and Hogan et al. also consider allocation between countries. Our analysis is differentiated by a deeper approach to the behavioral, demographic and decision models by addressing social distancing, essential worker groups, and allocation policies that can change over the course of the vaccination campaign.

General ethical guiding frameworks for vaccine prioritization decision-making have appeared earlier in the literature. Toner et al. (**5**) emphasize promoting three ethical values: the common good; fairness and equity; and legitimacy, trust and communal contributions to decision-making. Emanuel et al. (**4**) promote four ethical values: maximizing benefits, treating equally, instrumental value, and priority to the worst off. Our analytic focus on minimizing new infections, years of life lost (YLL), or deaths emerges from promoting "the common good" or "maximizing benefits". Our focus on essential worker groups illustrates how ethical values (e.g. prioritizing essential workers due to the fairness of protecting those placing themselves at risk) may overlap with the common good (e.g. prioritizing essential workers to best reduce mortality and transmission). Issues of fairness and equity and protecting the worst off are not directly analyzed here but remain critical considerations.

For the sake of simplicity, we do not address in detail the potential set of complex and differential feedback processes between health status and opening of schools, workplaces and other institutions. While we limit policy objectives to a concise metric of health outcomes (minimizing expected cases, YLL, or deaths) we acknowledge that other values of returning to school, work and social life are important. Finally, we do not address additional vaccine complications, such as temporary effectiveness, potential side effects or any failure to take a second dose of the vaccine if necessary.

Although much is known about the epidemiology of COVID-19, uncertainty remains a key limitation to modeling the disease. Therefore, we consider a wide range of plausible scenarios and focus on the general features of the solutions, the commonalities between the alternative scenarios, and identification of model parameters that drive systematic differences in optimal vaccine allocations. Given these assumptions we find that optimal allocation strategies are responsive to both the initial and evolving epidemiological landscape of the disease. When focusing on mortality (YLL or deaths), vaccination of older essential workers and ages 60+ was almost always a top priority (i.e. targeted in the first 30% of the population vaccinated). Alternatively, when infections are

minimized, essential workers are prioritized followed by school age children across a range of likely scenarios. We find that prioritization can substantially improve public health outcomes—31 to 40% in the Base scenario, relative to untargeted vaccination. Two components unique to our model are important contributors to this improvement. First, policies that differentiate and target essential workers in addition to age substantially outperform those utilizing age-alone. Furthermore essential worker differentiation reduces trade offs between objectives (e.g. deterioration of YLL and infection metrics when focused on minimizing deaths). Second, extending from a static allocation (without phases) to allowing changes in prioritization over time provides substantial gains. Finally, while optimal prioritization is quite insensitive to model specification when minimizing infections, we find some sensitivity when focused on minimizing deaths or YLL. This sensitivity indicates benefits to adjusting the targeting strategy at the local level to match epidemiological conditions.

2.4. Results

To illustrate the qualitative nature of optimal dynamic prioritization, we first present results from a single "Base" scenario, representing a plausible set of parameters (detailed in Appendix A.1.2). These results are then compared to a set of alternative model scenarios as described in Table 2.1. While we begin with Base scenario results, we emphasize the sensitivity analysis under alternative scenarios that follows since information about some input parameters—e.g. expected vaccination supply—continues to change with time. In Fig. 2.2 the Base model allocation decisions are shown for each monthly decision period (in percent of vaccine supply) and then cumulatively (in percent of group vaccinated) at three and six months, respectively. Broadly, we find that the optimal policy is very dynamic: specific groups are targeted each period and these targets shift over time. Furthermore, targeting is very narrow initially but then becomes less so as a larger fraction of the population has been covered.

The whiskers on bars in Fig. 2.2 show the range of alternative allocations that still produce an outcome that is within 0.5% of the optimum. These indicate that the optimized outcome is relatively sensitive to substitutions between groups for the first three months as indicated by narrow whiskers around the cumulative allocations. There is, however, some limited ability to substitute vaccines between the two essential worker groups in the first two months when minimizing YLL or deaths. As the size of the susceptible population declines due to vaccination and infections the optimized

outcomes become less sensitive to substitutions (longer whiskers) with shifts between nearly all groups possible without substantial sacrifice. This suggests that targeting strategies can become less strict over time as the most vulnerable populations are protected. Comparing individual periods (Fig. 2.2A) and cumulative measures (Fig. 2.2B) shows that whiskers represent a combination of substitution between groups as well as between periods for the same group.

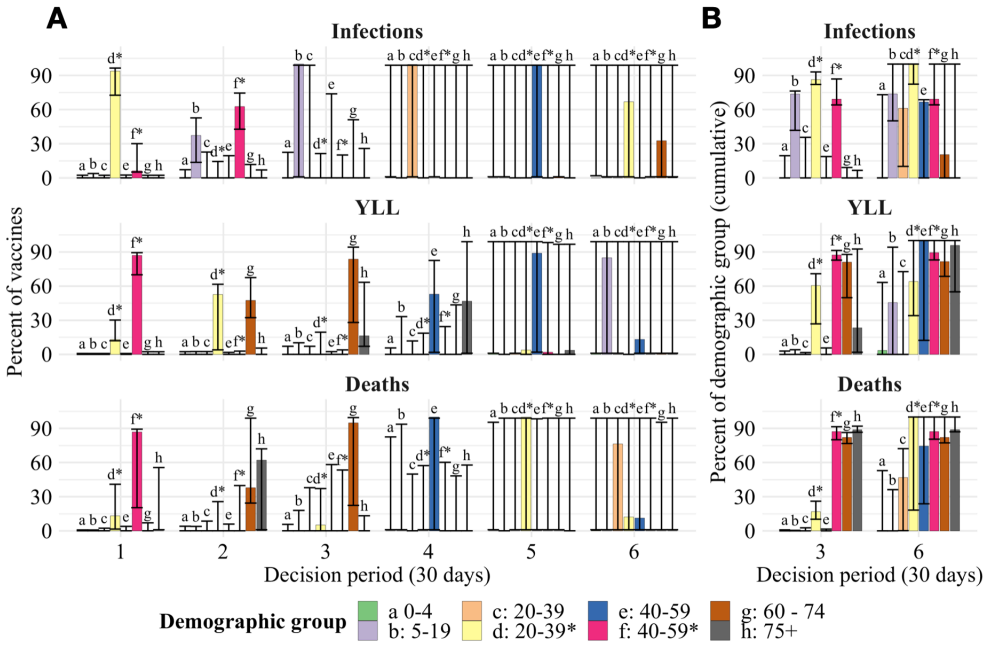

FIGURE 2.2. The optimal allocation of vaccines (vertical axes) between demographic groups for each decision period (horizontal axis) under the Base scenario (A). The three rows represent each objective, to minimize deaths, minimize years of life lost (YLL) and minimize infections. The bars for the six decision periods show the percentage of vaccines allocated to a specific group (indicated by a letter, color, and stars indicating essential worker groups) in that period. The two final columns (B) show cumulative measures at the end of months three and six, respectively, for the percent of each group that has been vaccinated. The whiskers on each bar represent the sensitivity of the optimal solution to small deviations in the outcome, specifically the range of allocations resulting in outcomes within 0.5% of the optimal solution.

Across objectives there are substantial differences in which groups are targeted early on. When minimizing deaths, targeting progresses from essential workers (20-39*, 40-59*), to the oldest (75+) and then younger seniors (60-74). These groups are a mix of those at high risk of mortality (older groups) and high risk of contraction and spread (essential workers). When minimizing YLL,

younger seniors are targeted earlier (given their longer average years of life remaining).[2] Finally, when minimizing infections we find that younger essential workers take top priority, followed by older essential workers and school-age children (5-19), since these groups have higher contacts and thus risk of contraction and spread.

In Fig.2.3A we show the dynamic path of infections, starting from the period in which vaccines become available, under various policies. As expected, infections are highest given no vaccines. Results for allocating vaccines in a manner proportional to each group's size shows the substantial value of even "untargeted" vaccines. As expected, the policy for minimizing infections leads to the lowest level of infections.

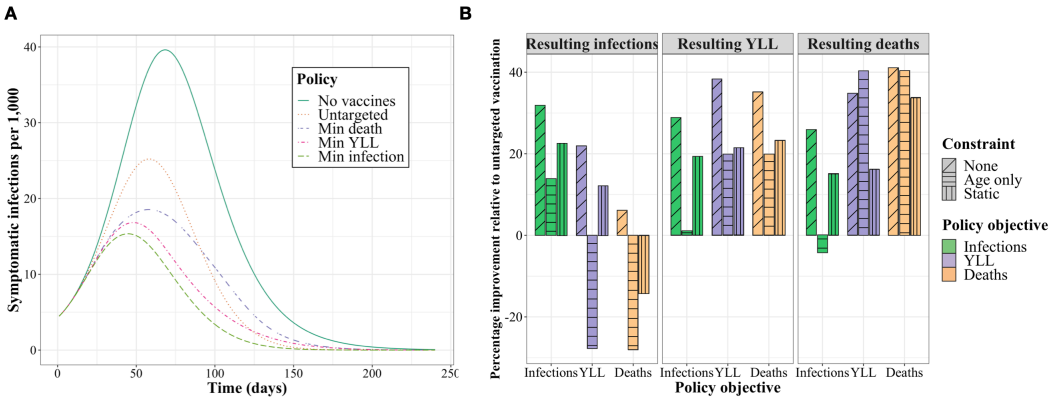

FIGURE 2.3. The number of infections per 1,000 individuals over time under reference policies (no vaccines; untargeted vaccine allocation) and optimized policies minimizing a given metric (A); and the performance of each optimized policy relative to an untargeted allocation policy (B) for the Base scenario. The bars are boxed by each resulting metric, colored by the objective driving each policy and textured to reflect any constraint considered (e.g. age-only or static policies).

In Fig. 2.3B we show the performance of various policies for resulting outcome metrics (infections, YLL and deaths) in terms of the percentage improvement relative to an untargeted vaccine allocation. We consider the optimal policies presented in Fig. 2.2 where the objective is minimizing infections (green), YLL (purple) , or deaths (orange) with no constraints ("none"). We also consider two constrained alternatives: an "age-only" dynamic policy that does not differentiate by essential worker status, and a "static" policy where the fractional allocation across groups does not change

[2]We do not discount in our calculation of YLL; doing so would lead to more equal weighting on mortality across age groups and thus results that are closer to those when minimizing deaths.

over decision periods.[3] We find that the unconstrained policy—that is dynamic and differentiated by essential workers—outperforms the untargeted approach by approximately 31-40% depending on the objective. Relative to the unconstrained policy, the age-only and static polices perform substantially worse for infections and YLL, though not for deaths. However, even while the age-only and static polices do not substantially impede performance in minimizing deaths, these constrained approaches still suffer substantial performance loss (9-18 percentage points) in the other two outcomes not optimized (YLL and infections) but clearly still of interest.[4] In other words, accounting both for essential workers and a dynamic prioritization strategy provides substantial improvements in the metric being optimized and/or the other two metrics of interest.

In general, we find that no matter the policy objective pursued in targeted vaccine allocation, some improvement is made on all three metrics. However, there are trade offs in what can be achieved between the objectives. For example policies that minimize infections result in substantially more deaths than a policy that minimizes deaths. We also find that differentiating essential workers substantially reduces these trade offs between objectives relative to age-only or static policies.

2.4.1. Sensitivity of vaccine prioritization. To assess how robust our Base scenario findings are to key uncertainties in the model, we conduct three different sensitivity analyses. First we consider a set of 10 alternative plausible scenarios involving a broad set of model inputs; then we focus on a narrower set of 4 parameters each explored in richer gradient detail; finally we examine a few fundamental changes to model structures.

2.4.1.1. *A broad set of alternative scenarios.* We solved for the optimal vaccine allocation across a range of 10 alternative scenarios selected to assess sensitivity to key assumptions of the Base model. Differences between these scenarios and the Base case are detailed in Table S.1. Relative to the Base model, in these alternative scenarios, we consider: higher initial infections; stronger or weaker non-social distancing non-pharmaceutical interventions (NPI) like mask wearing; weaker vaccine effectiveness overall or for seniors (60+); lower vaccine supply or supply that starts low and ramps up; more open schools; or higher contact rates overall.

To compare and contrast optimal early vaccination allocation for each scenario and objective, in Fig. 2.4 A we show the percentage of each group vaccinated after 30% of the overall population

[3]Excess vaccine is allocated without targeting if all the susceptible individuals in a given group have already been vaccinated.

[4]See resulting YLL and resulting infections for the deaths objective with age-only and static constraints.

Scenario	Change from Base scenario parameters	Source
Base scenario	None (Base parameter values are provided in Appendix A.1.2)	
High initial infections	Increased number of initial symptomatic infections (300% increase)	Assumed: pandemic state will vary between localities when vaccine first available to the general public.
Strong NPI	Non-social distancing (NSD) NPI are strong, resulting in a declining infection rate	Consistent with $R < 1$
Weak NPI	NSD NPI are weak, resulting in a sharply increasing burden of infection	Consistent with $R >> 1$
Weak vaccine	Lower vaccine effectiveness (success rate) for all age groups relative to the Base scenario	Minimum value required by FDA guidelines
Weak vaccine 60+	Lower vaccine effectiveness for ages 60+ yrs.	Informed by influenza vaccine effectiveness
Even susceptibility	All ages are equally susceptible to infection. Increase in susceptibility for ages < 20 yrs relative to Base	Assumed: tests sensitivity to age-dependent susceptibility described by (**9**) and (**19**)
Low supply	Sufficient supply for 5% of the population monthly (50% of supply relative to Base scenario; prioritization changes every 10% of the population vaccinated, such that decision period is 2 months)	Assumed: vaccine supply is uncertain and known to impact optimal allocations (**20**)
Ramp up	Vaccine supply is 5% per month for the first two months and 10% per month thereafter (first decision period is 2 months so increments of 10% of the population are vaccinated each decision period)	Informed by comments from the scientific head of the U.S. vaccine development program (**21**)
Open schools	Rate of social contact in schools increased from 30% in Base model to 70%	Assumed: tests sensitivity of optimal allocations to school closure intensity
High contacts	Increased number of contacts outside the home, school and workplace (50% increase relative to base)	Assumed: tests sensitivity to relaxed distancing

TABLE 2.1. Descriptions of alternative scenarios reletive to the Base model (see Appendix A.1.3 for specific levels).

is covered (typically in three months, except for alternative supply scenarios). We find that high priority groups—by percent of group vaccinated—are typically but not always robust to the alternative scenarios. For example, when deaths are considered (Fig.2.4 A, top panel) we see substitution between younger essential workers (20-39*) and ages 60-74 and when YLL are considered there is substitution between younger essential workers and ages 75+. To illustrate differences in the relative order of these high priority groups, in Appendix Fig. A.5 we show optimal prioritization of

vaccination in the very first decision period across objectives and scenarios. We find that when YLL are considered essential workers ages $40-59$ are the highest priority group in all scenarios. However, when deaths are considered ages 75+ are the highest priority group under several alternative scenarios.

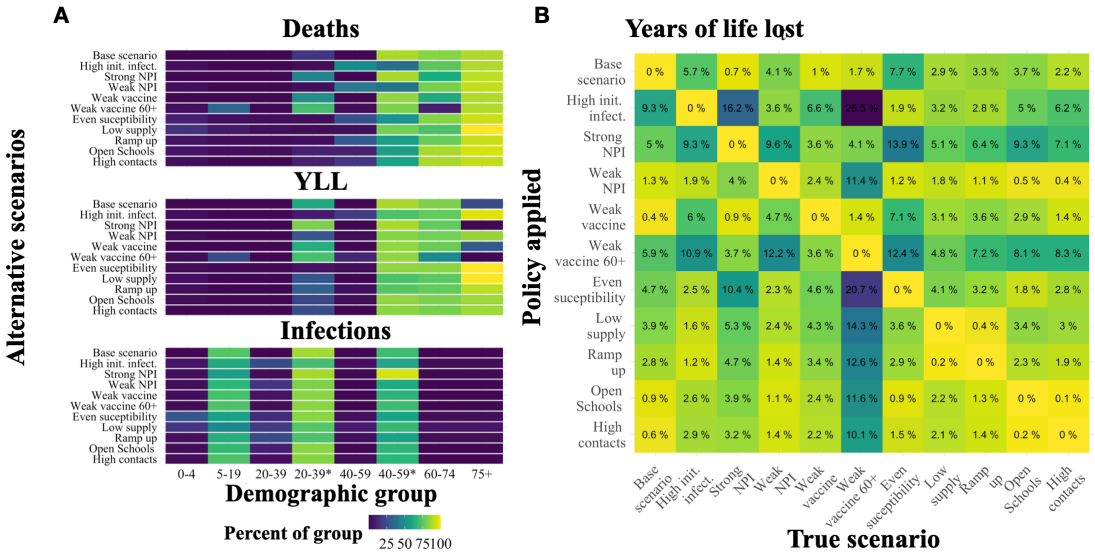

FIGURE 2.4. The cumulative percent of each demographic group (horizontal axis) vaccinated after the first 30% of the population is vaccinated under the alternative scenarios (vertical axis) and each objective (panel) (A). The percentage of additional YLL in excess of the optimum when applying a policy for a given alternative scenario (row) when a particular scenario is the "truth" (column) (B).

For insight into the cost of error in specifying the correct scenario, we assessed the performance of the policy identified for each of the 11 alternative scenarios, depending on which of these 11 is the "true" scenario. In Fig. 2.4B we show these results for the YLL objective. For example, the first column shows the performance loss (in percentage of additional YLL above the optimum) when the true scenario is the Base model but the decision maker applies a policy matched to any of the alternative scenarios (rows). By construction, when the policy applied matches the true scenario, the performance loss is zero. When YLL is the focus and the Base specification is the "true" scenario, the greatest performance loss (9%) comes from mistakenly applying the high initial infections policy.

We find that performance costs in percentage terms from applying the wrong policy from this set are typically modest (low single digits) albeit with notable exceptions. For example, when the "truth"

is that we have a weak vaccine for ages 60+, several policies applied perform very poorly relative to the true optimal policy since they substitute vaccine away from younger essential workers to ages 75+. A few of the policies were generally less robust across various true models, specifically those for high initial infections, strong NPI, and weak vaccine 60+. The Base scenario policy performed reasonably well across true alternative models, with the largest loss arising (7%) when children are not less susceptible (even susceptibility).

Equivalent versions of Fig.2.4B for minimizing deaths or infections are provided in the appendix (Figs. A.3, A.4). When the focus is minimizing deaths, the pattern of performance between scenarios is very consistent with YLL in Fig. 2.4. However, the scope for performance loss is larger overall–up from a maximum of 26% for YLL to 46% for deaths. When the focus is infections, the range of performance loss is much less intense at 7%. For infections, this relatively robust performance arises because optimal policies are much more similar across scenarios when minimizing infections (compared to the other objectives). Given greater scenario-driven heterogeneity in policies for minimizing YLL or deaths, there is greater opportunity for performance loss from specification error.

2.4.1.2. *A gradient over four key parameters.* For further sensitivity analysis, as shown in Fig. 2.5, we assessed how optimal vaccine allocation policy changed along a gradient for four key model inputs: non-social distancing NPI effectiveness (e.g. mask wearing) which determines the initial reproductive number (when the vaccine first becomes available); initial infections; monthly rate of vaccine supply; and vaccine effectiveness.

Echoing sensitivity results reported above, variation in these parameters had little effect on the optimal policy for minimizing infections. But we found systematic differences in the policies for minimizing YLL and deaths. Essential workers, ages 60-74 and ages 75+ remained the highest priority groups across the full range of parameters tested but there was substitution between younger essential workers (20-39*) and the older age groups.

In most instances the percent of vaccines responded in relatively monotonic fashion as parameters varied. For example, consider the objective of minimizing deaths. As depicted in Fig. 2.5A:D, prioritization of essential workers fell and 60+ or 75+ increased as (1) initial infections grow; (2) vaccine supply decreases; or (3) vaccine effectiveness increases. In a few instances, the percent of vaccine allocated to a given group responded non-monotonically to variation in the parameter. For

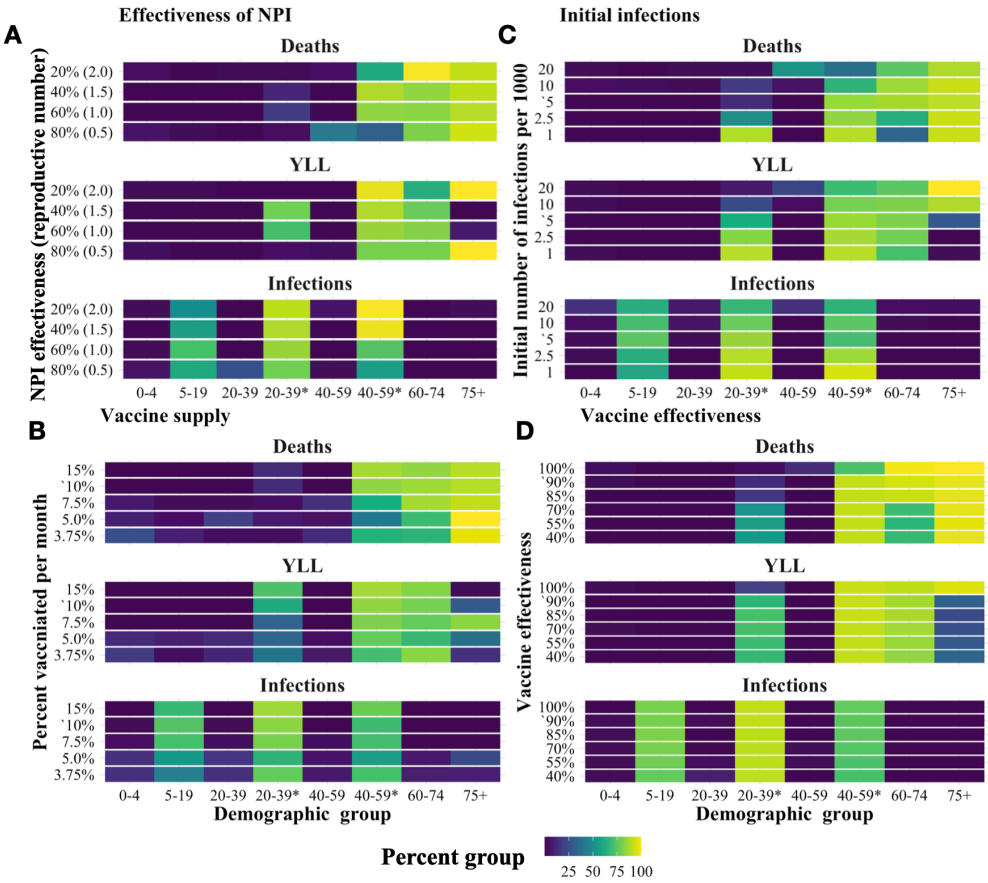

FIGURE 2.5. The total percent of each demographic group vaccinated after 3 months under the optimal dynamic policy. Each panel shows the effect of varying a key parameter relative to the Base model: (A) effectiveness of NPI, which determines the initial reproductive number (when the vaccine first becomes available); (B) monthly rate of vaccine supply; (C) initial infections; and (D) vaccine effectiveness. Base scenario parameter values are indicated with an apostrophe (‘).

example, for effectiveness of NPI in Fig. 2.5A, the allocation skewed towards 75+ and away from essential workers when the parameter was very high and very low.

These results indicate that when focusing on deaths or YLL, if transmission cannot be reduced quickly by the vaccine—due to limited supply, high reproductive numbers or large initial number of infections—typically this initial supply is most efficiently used to directly protect individuals with the greatest risk of death if infected. This pattern differs for vaccine effectiveness: we find that as the effectiveness of the vaccine decreased, supply is substituted away from the older (higher risk) age groups to essential workers. This difference is consistent with the fact that as vaccines become

less effective for a given individual, protecting vulnerable individuals is better achieved by reducing population-level transmission.

2.4.1.3. *Changes to model structures.* As a final sensitivity analysis, we examined robustness of the results to three alternative model structures: (1) clustered essential workers, where essential workers only contact other essential workers in the workplace, (2) concentrated essential workers, where relative to the Baseline scenario, the portion of the working age population deemed "essential" is half (20%) and they have approximately double the contact rate; and (3) leaky vaccine, where rather than working perfectly for 90% of individuals, vaccinated individuals have reduced susceptibility to infection, infectiousness and risk of death if infected. A more detailed discussion of these models is included in Appendix A.4.

We found that the qualitative nature of the solutions remained constant across each of these alternative models, with some minor differences. Treating the essential worker group as a cluster increased the proportion of vaccine allocated to ages 60+ when deaths and years of life lost were considered. This shows that when essential worker contacts are clustered within-group, this reduces the indirect protection that vaccinating these individuals provided to others. Conversely, concentrating the essential worker group (to a more select group with higher contact rates) increased the fraction of these individuals vaccinated. This shows that select essential workers with especially high contact rates (e.g. medical professionals and essential retail workers) are particularly strong candidates for early vaccination.

2.5. Discussion

Key insights and results from our analysis are summarized in Box 1. Together these lessons show the strong implications of considering dynamic solutions, social distancing and essential workers (given their limitations in social distancing) for vaccine prioritization. Our analysis of COVID-19 vaccine prioritization uniquely accounts for two critical needs: (1) dynamic prioritization given gradual roll out of vaccine during an active pandemic, and (2) attending to substantial heterogeneities in work contacts among the adult population due to the ability of many to work from home. These two novel features demonstrably change optimal vaccine prioritization. Given gradual vaccine deployment, static policies are out-performed by dynamic polices, which narrowly target a small number of demographic groups and (after substantial coverage of them) switch to lower priority

groups. Static policies identify a set of high priority groups but not how to order them when phased deployment is necessary. More strikingly, targeting essential workers (or other adults with large number of work contacts) reduces not just the adverse outcome of focus but also trade offs for remaining outcomes. For example, when minimizing deaths, allocation that differentiates essential workers substantially lessens the degree to which infections and YLL climb from the minimum achieved when each is optimized on its own.

Box 1. Key insights and results

(1) **Benefits:** Prioritization can reduce a particular undesirable outcomes (deaths, YLL, or infections), by 32-40% in the Base scenario (or 17-44% depending on the alternative scenario).

(2) **Objectives:** Moving from minimizing infections to YLL to deaths, boosts each of the following: benefits from vaccination targeting, prioritization differences between scenarios, and (therefore) the sensitivity of optimal prioritization to scenario.

(3) **Dynamic prioritization:** Dynamic prioritization (1) is responsive to the initial and evolving disease status, and (2) generates substantial improvement in outcomes relative to a static prioritization indicating that a phased approach to vaccine distribution is well justified. However, diminishing marginal returns to additional vaccination within a group drives a shift to other groups before 100% vaccination of the first group is achieved.

(4) **Widening prioritization:** As vaccination rates rise, precise prioritization becomes less critical and targeting widens to a larger set of groups.

(5) **Trade offs:** Policies that target one objective forgo opportunities to reduce alternative metrics. For example, policies that minimize deaths do not reduce infections nearly to the same degree as policies that minimize infections. These trade offs are typically stronger when policies do not allow for targeting based on essential worker status.

(6) **Essential workers:** Relative to an age-only model, policies that allow targeting of essential workers provide the greatest improvements when minimizing infections and YLL are the focus. In the Base scenario, essential workers are a high priority group under all three objectives (i.e. they are among the first 30% of the population to receive vaccines). However, their priority relative to ages 60+ is affected by key model parameters (see **Sensitivity** next).

(7) **Sensitivity:** The high priority groups remain consistent across the range of parameters considered. However, when minimizing deaths or YLL, the fraction of vaccine allocated to essential workers and ages 60+ depends on: the number of infections and reproductive number when the vaccine became available; the supply of vaccines; and vaccine effectiveness. In effect when the vaccine has a limited ability to quickly reduce the transmission of the virus optimal policies more heavily prioritize older individuals.

Existing published analysis of optimal COVID-19 vaccination targeting includes Matrajt et al. (**16**) and Bubar et al. (**17**). Before comparing and contrasting results some key modeling differences should be noted. These two analyses consider a wider range of vaccine availability than considered here. Our analysis uniquely incorporates non-pharmaceutical interventions (NPI), including social distancing and non-social distancing (e.g. mask wearing). Doing so allows us to account for differences between groups like essential workers constrained in distancing versus others who are much less so. All three preprints implement static optimization where vaccines are allocated and administered in a one-shot process. Our allocation is dynamic, responding to changing epidemiological conditions over a six-month period. Finally, all three model vaccines as "leaky", i.e., reducing the probability that a susceptible individual will be infected (and the probability of severe disease (**22**)). Bubar et al. also considers an "all-or-nothing" vaccine that is 100% effective for a fraction of the population. In our Base model the vaccine is "all-or-nothing", though we also consider a leaky vaccine, as discussed at the end of the Results.

Matrajt et al. (**16**) found that optimal strategies to minimize deaths and years of life lost will either exclusively target groups with high infection fatality rates maximizing the direct benefit of vaccines, or will target groups with high rates of infection maximizing the indirect benefits of the vaccine. In contrast, our results indicate that optimal policies initially target groups with high risk of infection and then switch to targeting groups with high infection fatality. This difference most likely follows from our dynamic versus static allocation. The switching behavior we identify is consistent with past work on pandemic influenza vaccine prioritization, which suggests that early in an outbreak when the infection rate is growing targeting spread (maximizing indirect benefits) is more efficient, but later when the infection rate is leveling off or declining, maximizing direct protection is most efficient (**20**).

Bubar et al. found that prioritizing adults older than 60 years of age is a robust strategy for minimizing deaths. In contrast we find that working-age adults are a key priority group, particularly older essential workers. These differences may either arise from our allowance for social distancing and/or dynamic allocation. Our accounting for social distancing on COVID-19 transmission increases the modeled benefits of targeting essential workers, who are less able to substantially reduce their social contacts than individuals over 60. Furthermore, as discussed above, the ability of dynamic policies to switch over time allows the allocation schemes we discuss to capture the benefits

of using the initial vaccine supply to slow transmission without sacrificing direct protection of more vulnerable individuals later on.

Two notable additional analyses of optimal COVID-19 vaccination targeting in preprint form include Wang et al. (**23**) who focus on the mortality costs of delay in vaccine rollout and the tradeoff between prioritizing first versus second doses as well as Hogan et al. (**22**) who examine ideal allocation both within and between countries.

National and international institutions have also begun to disseminate guidance. In particular, general guidelines for vaccine prioritization have been put forward by SAGE of the WHO (**24**) and the US CDC's ACIP (**25**). For example, CDC recommendations prioritize (A) health care personnel, (B) residents of long-term care facilities, (C) persons aged 75 and over and frontline essential workers, (D) persons aged 65-74, persons aged 16–64 years with high-risk medical conditions, and essential workers, and (E) everyone aged 16 and over remaining. (A) and (B) are subgroups at a finer scale than considered here, though with clear logic supporting top priority. Notably, a clear priority is not set between persons aged 75 and over and frontline essential workers. This is consistent with our findings in that prioritization within this pair was sensitive to specific conditions, which will vary over location. Our recommendations differ in the distinction made here between younger and older essential workers, with priority on the latter motivated by increasing mortality from infection with age. The CDC guidelines also consider underlying health conditions, a salient distinction not considered here. WHO guidelines—written more broadly for a global audience—agree on the prioritization of frontline health care workers at high risk of infection followed by older adults. However subsequent priority focuses on various sociodemographic groups at high risk (e.g. those in poverty) and essential educational workers before turning to essential workers more broadly.

Although our model provides useful insight for the policy-making process, a number of caveats are in order. In reality the risk of infection varies continuously across individuals, even between different "essential" occupations. While our model is unique in capturing differences between essential and non-essential workers, the representation of these differences is simplified by averaging the total number of contacts over a group with high work contacts (essential workers) and a group with lower rates of work contacts. This allows us to demonstrate the importance of this heterogeneity

in the adult population relative to the standard age-only models, indicating that policy makers should strongly consider occupation-differentiated vaccine allocation strategies.

While we explored a large set of alternative scenarios, further extensions remain for future work. For example, if certain groups (e.g., children or seniors) experience significant vaccination side effects, prioritization might shift away from these groups (**26**). From a logistical perspective, vaccination will occur through various points of contact with the community (pharmacies, clinics, schools, etc.). Constraints imposed by the distribution network used will affect the relative costs of reaching various subgroups. While the longevity of immunity to COVID-19—either following natural infection or vaccination—is not yet well understood, emerging analysis suggests that, following infection, "durable immunity against secondary COVID-19 disease is a possibility for most individuals" in the sense that immune memory was present in approximately 95% of individuals studied 5-8 months after symptom onset (**27**). How long lasting immune memory will be in the longer run is a key unknown. We assume immunity spans at least through the end of our six month time horizon. However, if instead this durability is more limited and/or already waning for those infected early in 2020, we might expect the symptomatic infections curve (Fig.2.3A) to stretch further out and for ideal vaccination strategies to shift towards direct protection of older, vulnerable populations.

From a behavioral perspective, vaccine hesitancy may influence the ability to achieve vaccination priorities, especially as coverage of the population increases. In general, we find that it is not necessary or even ideal to vaccinate all of the susceptible individuals in a demographic group, at least given the level of 60% of the population vaccinated considered here. Thus, at least initially, some level of vaccine hesitancy may have limited material impact. However, hesitancy may play a more significant role in the longer run, especially if hesitancy rates are large and herd immunity proves difficult to achieve (e.g. if vaccine effectiveness is low, and/or NPI relaxation is aggressive). Vaccine hesitancy that is concentrated in a particular community or demographic group could also change the optimal prioritization strategy. Similarly, adjustments would be needed if groups differ in the duration of vaccine effectiveness or diligence in obtaining a second dose of the vaccine where (and when) necessary.

For simplicity we limited policy objectives to a set of concise metrics of health outcomes (minimizing expected cases, YLL, or deaths). However, other health-related metrics such as protecting the most vulnerable and social values such as returning to school, work and social life are important to

consider. Our analysis reveals that optimal strategies for minimizing deaths and YLL are broadly aligned with the goal of protecting the most vulnerable. These solutions target essential workers who are the least able to participate in NPI such as social distancing and thus are the most at risk of infection, and individuals over the age of 60 yrs. who have the highest risk of death if infected by the disease. Other social values such as returning to school will most likely change the allocation schemes to offset the risk created by relaxing social distancing. For example, if allowing children to return to school was a high priority, then allocation strategies might be tilted towards targeting school-age children and teachers. A detailed analysis of optimal vaccine allocation given the relaxation of social distancing to achieve particular social objectives is a key direction for future research.

2.6. Methods

2.6.1. Model. To investigate the impact of vaccination strategies on the COVID-19 pandemic in the U.S., we employed a structured compartmental transmission model similar to (**28**). We incorporated the demographic structure of the population by tracking six age groups in the set $J = \{\text{0-4}, \text{5-19}, \text{20-39}, \text{40-59}, \text{60-4}, 75+\}$. We then extend this set to differentiate essential workers by splitting the two prime working age groups into two groups—non-essential workers (20-39, 40-59) and essential workers (20-39*, 40-59*)—yielding four groups of prime working age individuals and a total of eight demographic groups in $J = \{\text{0-4}, \text{5-19}, \text{20-39}, \text{20-39}^*, \text{40-59}, \text{40-59}^*, \text{60-74}, 75+\}$. For each demographic group we tracked 9 epidemiological states: susceptible (S), protected by a vaccine (P), vaccinated but unprotected (F), exposed (E), pre-symptomatic (I_{pre}), symptomatic (I_{sym}), asymptomatic (I_{asym}), recovered (R) and deceased (D). In Fig. 2.1 we display the compartmental diagram and directions of transitions between epidemiological states.

We modeled the COVID-19 transmission dynamics using a system of coupled ordinary differential equations for each demographic group, indexed by i and j. The transmission rate was given by the product of the transmission probability (q), the age-specific susceptibility (s_i), strength of non-pharmaceutical interventions (θ), the relative infectiousness of each symptom type (τ_m)—where $m \in M \equiv \{asym, pre, sym\}$—and the rate of contact ($r_{m,i,j}$) between infected individuals with symptom type m from group j and susceptible individuals from group i. The exogenously given

population vaccination rate at time t is given by $v(t)$, where units of time are days.[5] In our Base model we assume that for each individual the vaccine either works or it does not (though we also consider vaccines that are partially effective for all vaccinated in our sensitivity analysis). Individuals in group i are vaccinated at a rate of $\mu_i v(t)$ and a fraction of the those (ϵ_i) are protected while a fraction remain susceptible and move to the failed vaccination category (F).[6] Once infected, individuals move from exposed to pre-symptomatic at rate γ_{exp}^{-1}. Pre-symptomatic individuals become symptomatic or asymptomatic at rates $\sigma_{asym}/\gamma_{pre}$ and $(1-\sigma_{asym})/\gamma_{pre}$, respectively. Asymptomatic individuals recover at an uniform rate γ_{asym}^{-1} and symptomatic individuals either recover or die at a rate of $(1-\delta_a)/\gamma_{sym}$ or δ_a/γ_{sym}, respectively, where δ_a is the age-specific infection fatality rate. These assumptions yield the system of differential equations described in Appendix A.1.1, with parameter values given in Appendices A.1.2 and A.1.3.

2.6.2. Contact rates. Contact rates indicating the level of direct interaction of individuals within and between groups drive the transmission dynamics in the model. We built the contact matrices used in this model from the contact matrices estimated for the U.S. in (**14**). These estimates are given for age groups with five year age increments from 0 to 80 yrs. These estimates were aggregated to provide estimates for the six-level age structure used in our model. We also extended these data to estimate the contact rates of essential workers. A detailed derivation of these contact rates can be found in Appendix A.1.6. In short, we assumed that essential workers have on average the same pattern of contacts as an average worker in the population in the absence of social distancing. We then scaled the contact rates for essential and non-essential workers to represent the effects of social distancing and calculated the resulting mixing patterns assuming homogeneity between these groups.

We constructed contact matrices for each of four locations, $x \in \{home, work, school, other\}$, following (**14**). The total contact rate for an asymptomatic individual before the onset of the pandemic is given by the sum of these location specific matrices. However, it is clear that populations are exhibiting social distancing in response to the pandemic (**29**). We further expect symptomatic

[5]In the event that vaccination requires two doses over time, we consider an individual vaccinated upon receipt of the second dose at time t and we assume that v indicates the number of *individuals* that can be vaccinated with the required number of doses.

[6]This vaccine effectiveness is inclusive of any efficiency loss from typical handling in the distribution chain.

individuals to change their behavior in response to the illness. We account for these behavioral changes as described next.

2.6.3. Social distancing. Expression of symptoms and social distancing policies are likely to change individuals' behaviors over time. To model these changes we scaled the contribution of each contact matrix for location x:

$$r_m = \sum_x \alpha_{m,x} * r_x. \tag{2.1}$$

The weights $\alpha_{m,x}$ depend on disease and symptom status (m) and location (x) as specified in Table 2.2. We scaled social contacts for symptomatic individuals following changes in behavior observed among symptomatic individuals during the 2009 A/H1N1 pandemic (**30**). For those without symptoms (susceptible and asymptomatic) the weights were specified to match reduced levels of social contacts as the product of social distancing policies. Home contact rates were held constant, and non-household contact rates were roughly based on survey data from (**15**). However, levels of social distancing have varied strongly over time and between locations. To account for this variability we tested a range of alternative levels in addition to the Base model. The results for these alternative parameter values are discussed in Appendix A.1.6. Also, notably we do not consider the seasonality of contact rates for children in the scenarios where schools are modeled as closed. This would likely have limited impact on the optimal solutions, but when this is not the case we may over or underestimate the importance of school contacts depending on the time of year when vaccines are distributed.

The proportion of essential workers in the population was set to be consistent with estimates of the portion of jobs that can be done from home (**31**) and estimates from the U.S. Cybersecurity and Infrastructure Security Agency, which indicate that 70% of the workforce is involved in these essential activities (e.g. healthcare, telecommunications, information technology systems, defense, food and agriculture, transportation and logistics, energy, water, public works and public safety) (**32**). However, essential workers are not a cleanly defined group of individuals and there is heterogeneity in the level of contact rates within this group. As a robustness check on this Base scenario approach, we also tested a model with a smaller number of essential workers with higher contact rates. Results from this model are discussed in Appendix A.4.

Disease and symptom type	Contact rate weights, $\alpha_{m,x}$			
	home	*work*	*school*	*other*
symptomatic	1.0	0.036	0.036	0.075
susceptible or asymptomatic	1.0	0.4*, 0.1	0.3	0.4

TABLE 2.2. Weights on contact rates for a given disease and symptom type (m) and location/activity (x) under social distancing. When essential and non-essential-worker weights are both needed the former is marked with a star.

2.6.4. Transmission rate and vaccine effectiveness. The model was calibrated to match the predicted R_0 for COVID-19 in the U.S. (see Appendix Table A.1.2) by solving for probability of transmission q, assuming a naive (pre-pandemic) population. Details of this procedure are provided in Appendix A.1.5.

In our Base model we considered vaccine effectiveness of 90%. This level is at the low end of the range of estimates reported (90-95%) for reduction in symptomatic infections in the fall of 2020 from phase three clinical trials (**33**). We selected the low end since real world performance is typically somewhat lower than clinical trial effectiveness due to imperfect implementation of dual-dose timetables and or cold storage requirements. We also assume this effectiveness is the same across age groups since initial evidence does not show substantial differences between subgroups (**34**). As an alternative, lower-bound scenario we considered vaccine effectiveness of 50% since this is the minimum expectation of the U.S. FDA for approval (**35**). Finally, we considered a case where the vaccine is less effective for ages 60+. The phase three trials do not fully resolve the effectiveness of the vaccines by age, leading to uncertainty. This scenario represents a worst case scenario where the vaccine is much less efficacious for the most sensitive groups.

2.6.5. Initial conditions. Because the expected epidemiological conditions $\{I_{pre}(0), I_{asym}(0), I_{sym}(0), S(0)\}$ by the time the initial vaccine doses are ready for deployment are uncertain, we consider a range of possible values from 1 case per thousand to 20 cases per thousand. These cases were apportioned between demographic groups to reflect the attack rates of COVID-19 for each group under the given social distancing policy. Alternative levels considered for initial conditions are described in Appendix A.1.4 and appear in Results (see Table 2.1, Fig. 2.4 and Fig. 2.5C).

2.6.6. Vaccine prioritization optimization. The planner's decision problem is to allocate the daily supply of vaccine ($v(t)$) across the demographic groups according to a given objective.

We assume that this group allocation vector, μ, can be chosen on a monthly basis at the beginning of each of the first six decision periods. We also assume that only susceptible individuals are vaccinated. We numerically solved for vaccine allocation strategies that minimize the total burden associated with three different health metrics: deaths, years of life lost (YLL) or symptomatic infections:

$$\text{deaths: } \min \left\{ \int_0^T \sum_{i \in J} I_{sym,i}(t)/\gamma_{sym} dt \right\} \tag{2.2}$$

$$\text{YLL: } \min \left\{ \int_0^T \sum_{i \in J} e_i \delta_i I_{sym,i}(t)/\gamma_{sym} dt \right\} \tag{2.3}$$

$$\text{symptomatic infections: } \min \left\{ \int_0^T \sum_{i \in J} \delta_i I_{sym,i}(t)/\gamma_{sym} dt \right\}, \tag{2.4}$$

where e_i is the years remaining of life expectancy for group i and with a six month time horizon ($T =$ 180 days). Preventing deaths and years of life lost are "consensus value(s) across expert reports" (**4**, p. 2052) while some argue that "protecting public health during the COVID-19 pandemic requires...minimizing COVID-19 infection" (**5**, p. 10).

We solved for the optimal allocation of available vaccines across demographic groups for each month over six months. We identified the optimal solution using a two-step algorithm. In the first step we used a genetic algorithm similar to (**36**) to identify an approximate solution. This approach uses random sampling of the potential solution space to broadly explore in order to avoid narrowing to a local and not global minimum. In the second step we used simulated annealing to identify the solution with precision. At a given optimal solution, it may or may not be the case that the outcome of interest (e.g. minimizing deaths) is sensitive to small changes in the allocation decision. Thus, around the optimal allocation we also identified nearby allocations that produce outcomes that are less desirable but still within 0.5% of the optimized outcome. A detailed description of the algorithm is given in Appendix A.1.7. All code for the optimization was written in the Julia programming language (**37**).

The whiskers on optimal vaccine allocation bars in Fig. 2.2 show the range of alternative allocations that still produce an outcome that is within 0.5% of the optimum. The upper (lower) bound of each whisker was produced one at a time by systematically exploring higher (lower) levels of the given decision variable (proportion of vaccines allocated to a given demographic group in a given

decision period). This entailed fixing a candidate level of that decision variable as a constraint and optimizing the remaining parameters. If the optimized value of this constrained objective function was within 0.5% of the unconstrained optimum then the candidate value was accepted and included within the bounds specified by the whisker. The whisker bounds were found using bisection linear search algorithm tuned to identify each bound to within one percentage point of the true value (see Appendix A.1.8).

To assess the benefits of (1) using a dynamic allocation policy and (2) differentiating by essential worker status in addition to age, we constructed two constrained policies: a static policy and an age-only policy. The static policy was found by allowing the proportion of vaccine allocated to each age group to be chosen once when the vaccine first becomes available and then applied constantly over time.[7] The age-only policy simply involves constraining allocation choices age groups (not differentiated by essential worker status)—vaccines allocated to working age groups accrue to essential workers simply in proportion to their relative share of these groups.

[7]If all of the susceptibles from a single group were exhausted (either by full coverage from the vaccine or from infection) then vaccine that would have been allocated to individuals from that group are instead allocated to other age groups at a rate proportional to the size of their susceptible population.

Bibliography

1. L. Corey, J. R. Mascola, A. S. Fauci, F. S. Collins, *Science* **368**, 948 (2020).
2. E. Dong, H. Du, L. Gardner, *The Lancet Infectious Diseases* **20**, 533 (2020).
3. N. Florko, First Covid-19 vaccines to arrive in states Monday, marking a pivotal moment in the pandemic response, STAT, https://www.statnews.com/2020/12/12/shipments-of-first-authorized-covid-19-vaccine-in-the-u-s-will-begin-monday/ (2020).
4. E. Emanuel, *et al.*, *The New England Journal of Medicine* **382**, 2049 (2020).
5. E. Toner, A. Barnill, C. Krubiner, et al., Interim framework for COVID-19 vaccine allocation and distribution in the United States, Baltimore, MD: Johns Hopkins Center for Health Security, https://www.centerforhealthsecurity.org/our-work/pubs_archive/pubs-pdfs/2020/200819-vaccine-allocation.pdf (2020).
6. M. Lipsitch, N. E. Dean, *Science* **370**, 763 (2020).
7. N. Lurie, M. Saville, R. Hatchett, J. Halton, *New England Journal of Medicine* **382**, 1969 (2020).
8. J. Medlock, A. P. Galvani, *Science* **325**, 1705 (2009).
9. N. G. Davies, *et al.*, *Nature Medicine* **26**, 1205 (2020).
10. R. M. Viner, *et al.*, *JAMA Pediatrics* (2020).
11. R. Verity, *et al.*, *The Lancet Infectious Diseases* **20**, 669 (2020).
12. M. C. Fitzpatrick, A. P. Galvani, *Science* (2021).
13. J. Mossong, *et al.*, *PLoS Medicine* **5**, e74 (2008).
14. K. Prem, A. R. Cook, M. Jit, *PLoS Computational Biology* **13**, e1005697 (2017).
15. D. Feehan, A. Mahmud, Quantifying population contact patterns in the United States during the COVID-19 pandemic, *preprint*, Epidemiology (2020).
16. L. Matrajt, J. Eaton, T. Leung, E. R. Brown, *Science Advances* **7** (2021).
17. K. M. Bubar, *et al.*, *Science* **(online early)** (2021).
18. M. E. Gallagher, *et al.*, Considering indirect benefits is critical when evaluating sars-cov-2 vaccine candidates, *preprint*, medRxiv, doi: 10.1101/2020.08.07.20170456 (2020).
19. J. Zhang, *et al.*, *Science* **368**, 1481 (2020).
20. L. Matrajt, M. E. Halloran, I. M. Longini Jr, *PLoS Comput Biol* **9**, e1002964 (2013).

21. M. Slaoui, Interview with Mary Louise Kelly: Operation Warp Speed top adviser on the status of a coronavirus vaccine, National Public Radio, https://www.npr.org/2020/09/03/909312697/operation-warp-speed-top-adviser-on-the-status-of-a-coronavirus-vaccine (2020).
22. A. Hogan, *et al.*, Modelling the allocation and impact of a COVID-19 vaccine, *preprint*, Imperial College London (2020).
23. X. Wang, *et al.*, The impacts of COVID-19 vaccine timing, number of doses, and risk prioritization on mortality in the US, *preprint*, medRxiv, doi: 10.1101/2021.01.18.21250071 (2021).
24. S. Omer, *et al.*, WHO SAGE roadmap for prioritizing uses of COVID-19 vaccines in the context of limited supply, World Health Organization, https://www.who.int/publications/m/item/who-sage-roadmap-for-prioritizing-uses-of-covid-19-vaccines-in-the-context-of-limited-supply (2020).
25. K. Dooling, *et al.*, *Morbidity and Mortality Weekly Report* **69**, 1657 (2021).
26. National Academies of Sciences, Engineering, and Medicine, Discussion draft of the preliminary framework for equitable allocation of COVID-19 vaccine, Washington, DC: The National Academies Press. https://doi.org/10.17226/25914 (2020).
27. J. M. Dan, *et al.*, *Science* **371** (2021).
28. S. Abrams, *et al.*, Modeling the early phase of the Belgian COVID-19 epidemic using a stochastic compartmental model and studying its implied future trajectories, *preprint*, medRxiv, doi: 10.1101/2020.06.29.20142851 (2020).
29. J. A. Weill, M. Stigler, O. Deschenes, M. R. Springborn, *Proceedings of the National Academy of Sciences* **202009412** (2020).
30. K. Van Kerckhove, N. Hens, W. J. Edmunds, K. T. Eames, *American Journal of Epidemiology* **178**, 1655 (2013).
31. A. W. Bartik, Z. B. Cullen, E. L. Glaeser, M. Luca, C. T. Stanton, What jobs are being done at home during the COVID-19 crisis? Evidence from firm-level surveys, *National Bureau of Economic Research Working Paper No. 27422* (2020).
32. M. Romitti, Many U.S. workers in critical occupations in the fight against COVID-19, Labor and Market Information Institute, https://www.lmiontheweb.org/more-than-half-of-u-s-workers-in-critical-occupations-in-the-fight-against-covid-19/#:~:text=According%20to%20the%20federal%20standard,Workforce%E2%80%9D%20battling%20COVID%2D19 (2020).
33. M. Herper, D. Garde, Moderna to submit Covid-19 vaccine to FDA as full results show 94% efficacy, STAT, https://www.statnews.com/2020/11/30/moderna-covid-19-vaccine-full-results/ (2020).
34. C. Zimmer, Moderna's Covid vaccine: What you need to know, New York Times, https://www.nytimes.com/live/2020/moderna-covid-19-vaccine (2020).
35. U.S. Food and Drug Administration, Development and licensure of vaccines to prevent COVID-19: Guidance for industry, https://www.fda.gov/regulatory-information/search-fda-guidance-documents/development-and-licensure-vaccines-prevent-covid-19 (2020).

36. R. Patel, I. M. Longini Jr, M. E. Halloran, *Journal of Theoretical Biology* **234**, 201 (2005).

37. J. Bezanson, A. Edelman, S. Karpinski, V. B. Shah, *SIAM Review* **59**, 65 (2017).

CHAPTER 3

Long life spans can mitigate the genetic effects of strays from temporary conservation hatchery programs.

3.1. Abstract

While artificial propagation can serve as a conservation tool that provides demographic support, it can also cause domestication and outbreeding depression that risks unintended fitness consequences to both target and connected populations. When artificial propagation programs are used temporarily, the interaction between program duration and species life history might determine the potential for such fitness consequences and the effectiveness of mitigation measures. Here, we develop a mathematical model to quantify approaches to mitigate unintended fitness consequences in temporary artificial propagation programs. We build and parameterize our model on a conservation aquaculture-based recovery program for white sturgeon (Acipenser transmontatus) in the Nechako River (British Columbia), which might impose genetic risks on the adjacent populations within the Fraser River. We find that, within time scales of 50-200 years, the life history characteristics of white sturgeon, particularly late age of maturity and longevity, typically reduce the genetic risks associated with captive breeding, compared to shorter-lived species like salmonids. Genetic effects of gene flow from hatchery-origin individuals to nearby wild populations accumulate at a rate inversely proportional to the generation time of the population, and thus slowly in long-lived species. The slow rate at which risks accumulate in this system provides a window of opportunity to learn and adapt management while the risks of hatchery inputs remain relatively small.

3.2. Introduction

Captive breeding and artificial propagation can be valuable conservation and natural resource management tools that help achieve a range of objectives. For example, fish hatcheries are used to supplement harvested populations (**1**,**2**), maintain declining populations of threatened and endangered species (**1**, **3**), and re-introduce locally extirpated species into formerly occupied habitats (**3**, **4**). These practices, however, can have unintended consequences that threaten biodiversity through genetic and ecological mechanisms (**5**). Captive breeding can affect the fitness of supplemented populations through inadvertent domestication (**6**, **7**) and reduce genetic diversity by increasing the reproductive output of a few individuals (**8**, **9**). Individuals released from captive breeding programs can also have unintended ecological interactions with their wild counterparts through competition for scarce resources, increasing predator abundance, and the risk of disease (**10**).

In addition to these potential impacts on the target population, hatchery production can impact neighboring wild populations when hatchery-origin fish are introduced to open systems and "stray" between adjacent locations. This dynamic creates the possibility that hatcheries designed to support a single population can threaten the viability of neighboring groups and the collective meta-population (**1**,**11**). In addition to the risks from domestication, genetic diversity, and ecological interactions analogous to the risks to the target population, straying poses additional risks to non-target populations. First, outbreeding depression can occur if straying leads to hybridization between populations with distinct local adaptation, thus reducing fitness (**1**). Second, straying can also threaten diversity by homogenizing formerly distinct populations, affecting population dynamics by increasing synchrony and reducing portfolio effects (**12**,**13**).

Previous modeling studies suggest that even a small number of hatchery strays can have large effects on the long-term genetic and demographic state of wild populations (**14**,**15**,**16**). For example, Ford (2002) showed that the equilibrium genetic state of the wild population can approach that of the hatchery population unless the rate of gene flow is very small (less than 5% of the spawning population is of hatchery origin). The models used to derive these results, however, describe the long-run equilibrium behavior of the system. As a result, these model findings might not hold if hatchery production is intended as a temporary conservation measure, which is often the case due to its high costs (**17**) and use as a stop-gap measure until other measures address the root cause of the species decline (**1**,**18**). Furthermore, declines in fitness and genetic diversity caused by captive breeding can occur more slowly in longer-lived species (**19**). Therefore, life history and duration of stocking could interact to mitigate the potential unintended consequences of hatchery production. Species life history and the duration or hatchery programs can also determine the effectiveness of management strategies designed to mitigate the genetic effects of strays. Potential measures to reduce the impacts of the hatchery releases include reducing the number of individuals released, removing individuals that have strayed to non-target populations, and limiting the duration of the stocking program. The relative efficacy of these approaches, however, might depend on their interactions with one another. For example, the effect of reducing stray rates could depend on the duration of the stocking program because a longer duration would allow more time for fitness effects to accumulate in both the hatchery-raised and wild-origin populations. Furthermore, the effectiveness of removing hatchery-origin individuals that have strayed from adjacent populations

could be determined by their life history because reproductive output and thus potential for gene flow change though an individual's life cycle.

An example of a temporary conservation hatchery with potential effects on connected populations is the conservation fish culture program for Nechako white sturgeon (Acipenser transmontanus). White sturgeon are one of 27 species of sturgeon (Ascipenseridae), freshwater and anadromous fishes native to the Holarctic. Across this range, all species have been assessed at some level of risk of extinction and all are characterized by a long-lived life history. Populations in the upper Fraser River, British Columbia, form a genetically distinct designable unit (DU) have been assessed as Endangered by the Committee on the Status of Endangered Wildlife in Canada (COSEWIC), and the DU awaits a listing decision under the Canadian Species at Risk Act (**20**). The Upper Fraser DU contains three groups that show a combination of genetic, demographic, and spatial distinctions the Nechako River, mid-Fraser, and upper-Fraser (**21**, **22**, **23**). The two groups in the mainstem of the Fraser River (mid-Fraser and upper-Fraser) are near historical abundance levels (**24**). The Nechako River group, however, has experienced a recruitment failure and a long-term decline in abundance (**25**, **26**).

To forestall further effects of the ongoing recruitment failure a captive breeding program was established to supplement recruitment while research and restoration efforts address the root causes of the decline. Hatchery production began in the Nechako River in 2015 with the goal of releasing a sufficient number of juveniles to produce 50-60 mature adults per year class over a 45-year period (**27**). Since hatchery releases began, evidence has accumulated that hatchery-origin fish have emigrated from the Nechako River to the mid-Fraser, either permanently or temporarily. In some areas, hatchery fish comprise over half of the juveniles captured in monitoring programs in the mid-Fraser River (**28**). This unexpectedly high proportion of hatchery-origin fish has raised concerns that the Nechako River hatchery might impact the neighboring populations in the mainstem of the Fraser River. Since inception, the captive rearing program has been designed to minimize the genetic risks by using wild-origin brood stock from the Nechako River and maximizing genetic diversity by using full factorial crosses. These steps, however, do not rule out the possibility that hatchery-origin fish that stray to the Fraser River could cause outbreeding depression. Declines in fitness could occur if the Nechako-origin fish dilute local adaptation in the mid-Fraser population or if they experience rapid adaptation to captivity, as has been documented in salmonids (**6**, **7**).

The long-lived life history of white sturgeon and the intention to use the hatchery as a temporary conservation measure, however, might reduce genetic and ecological risks.

To explore how species life history and hatchery management measures might affect the potential for unintended fitness consequences of temporary hatchery programs, we developed a mathematical model of hatchery immigrating into a wild population. To build and parametrize our model, we use the Nechako River hatchery and neighboring population segments in the mainstem Fraser River as an example system. Building on previous work (**14**, **16**, **29**) we use a coupled genetic-demographic model to describe the effects of the hatchery program on the abundance and fitness of neighboring wild populations. To capture the effects of the population's life history, we use an age-structured demographic model parameterized to match demographic rates of white sturgeon measured by a system-wide monitoring program (**21**). In addition, we develop a novel method to calibrate the genetic parameters of the model to match fitness declines observed in early-generation hatchery fish estimated by meta-analysis (**30**), which provides additional empirical grounding for our modeling framework.

Using this modeling framework, we test how interactions between the duration of the hatchery program and the species' longevity interact to determine the fitness and demographic effect of the hatchery program, while accounting for feedbacks between these demographic and genetic changes. We also test how life history and transient dynamics of temporary hatcheries determine the effectiveness of alternative mitigation strategies and discuss implications of our findings for adaptive management.

3.3. Methods

3.3.1. Model overview. Our model describes the genetic and demographic consequences of hatchery-origin immigrants on a neighboring wild population by tracking the abundance and genotypes of the wild population as it receives an influx of immigrants from a hatchery-raised population. We develop models for both the Nechako River hatchery population and a single wild population to represent the naturally occurring Fraser River populations. Given the recruitment failure in the Nechako River and the corresponding lack of natural production, we assume all reproduction in the system occurs in captivity and is subject to domestication selection. Following previous modeling studies (**14**, **16**), we represent the effects of domestication with a single quantitative genetic trait.

This trait represents a suit of phenotypes that can adapt to the natural or hatchery environment. We model domestication selection with stabilizing selection acting on the fecundity of adults and the viability of yearlings. We calibrate the strength of domestication selection to match fitness declines observed in early-generation hatchery fish (**30**). A fixed reduction in the fecundity of hatchery-origin individuals in the wild represents the plastic effects of captive breeding on fitness. To capture the effects of the domestication of hatchery origin stray on the Fraser population, we track the abundance and genotype distributions through five demographic events: reproduction, density dependence, selection, immigration, and adult survival (Figure 3.1). For reproduction, fecundity is a function of age and genotype to capture the effects of growth, maturation, and fecundity selection. Mating pairs form independently of their ages and phenotypes, and using the infinitesimal model of quantitative genetics, the genotypes of their offspring follow a normal distribution centered at the mid-parental value (**31**). Yearlings experience density-dependent mortality and viability selection around an optimum genotype, with selection events occurring before and after density-dependence to represent a combination of hard and soft selection, respectively (**16**,**32**). We then track each cohort as it ages and experiences growth, maturation, natural mortality, and senescence. We parameterized the vital rates of the population to match observations from a long-term monitoring study (**21**) and previous demographic analysis (**24**, **26**).

To understand how hatchery operations affect the fitness and abundance of the wild population, we varied three factors: (1) the number of hatchery-origin individuals that immigrated into the wild population per year (immigration rate), (2) duration of hatchery releases, and (3) stray removal. The immigration rate captures the effect of the number of individuals released from the hatchery and factors influencing the probability of straying between river segments, such as acclimatization to the local environment before release. To test the effect of the duration of hatchery production and the system's dynamics, we vary the time over which immigration from the hatchery population occurs in our numerical simulations. For stray removal, we test the effect of removing post-release hatchery strays of different age classes and how the effectiveness of removals changed as a function of the planned hatchery duration and immigration rate.

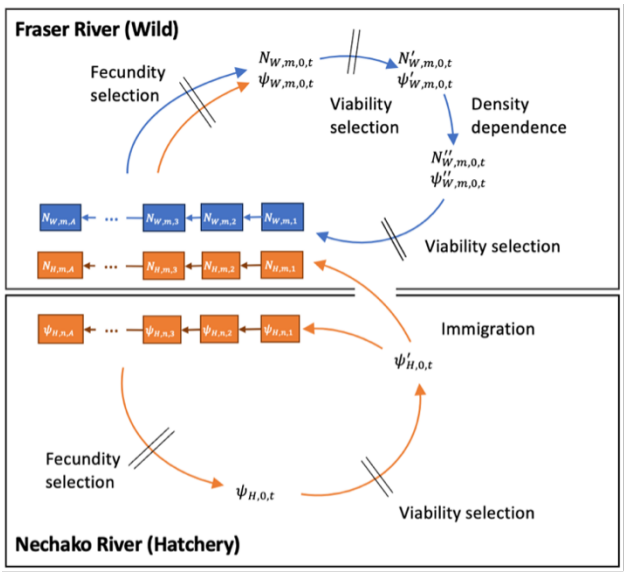

FIGURE 3.1. The model is composed of two components: a model of the genotypes of the Nechako River population as it experiences domestication selection and a model of the fitness and abundance of the wild Fraser River population as is receives immigrants from the Nechako hatchery population. In both models, fecundity selection affects the reproductive success of adults, and viability selection affects the survival of juveniles. In addition, the model of the Fraser River population includes density-dependent mortality. Viability selection occurs before and after density dependence, representing soft and hard selection. Immigration from the Nechako River to the mainstem Fraser occurs before recruitment at age one, and immigrants remain in the mainstem or Nechako River population for the remainder of their life cycles. Each of these events effects the abundance $N_{(i,j,a,t}$ and genotype distribution $\psi_{i,j,a,t}$of the population of origin i in location j. The arrows indicate moment between demographic stages. Transitions influence by selection are indicated by hash marks.

3.3.2. model details.

3.3.2.1. *State variables.* We follow the population size of each cohort $N_{i,j,a,t}$ and their genotype g distribution $\psi_{i,a,t}(g)$ where i denotes the origin (hatchery H or wild W), j denotes their location (Nechako River n or Main-stem Fraser River m), a denotes age, and t denotes time. We follow the full genotype distribution rather than tracking the mean and variance of a normal distribution because gene flow between the populations can create departures from normality (**16**). We track the abundance of wild origin yearlings $N_{W,m,0,t}$ through three life stages prior to recruitment, with each stage denoted by $N_{W,m,0,t}$, $N'_{W,m,0,t}$,, and, $N''_{W,m,0,t}$, along with the genotype distributions $\psi_{W,m,0,t}(g)$, $\psi'_{W,m,0,t}(g)$, and, $\psi''_{W,m,0,t}(g)$. We track the genotypes of hatchery origin fish through two

life stages prior to recruitment: $\psi_{H,n,0,t}(g)$ and $\psi'_{H,n,0,t}(g)$. We only track the genotype distribution in the hatchery environment assuming a management-controlled constant population size.

3.3.2.2. *Reproduction.* The fecundity of an individuals is determined by their age and genotype, and environmental effects of the hatchery, where the effect of age is given by a vector of age-specific fecundities F_a, and the effect of genotypes is given by a selection gradient $W_{f,j}(g)$. Fecundity selection is stabilizing around an optimum genotype θ_j with strength $s_{f,j}$, both of which depend on the spawning environment $j \in \{W, H\}$. Regardless of genotype, Hatchery fish spawning in the wild are less successful by a factor of Δ than their wild origin counterparts. Therefore, the fecundity of an individual in spawning environment j is the product of these three factors

$$F_{j,i,a}(g) = F_a \Delta^{I_{i=H}} e^{-0.5 s_{f,j}(g-\theta_j)^2} \tag{3.1}$$

where $I_{i=H}$ is an indicator variable that takes the value of one for hatchery fish spawning in the wild and are zero otherwise. Following Wood et al. (2007) and Smyth et al. (2016), fecundity at age is proportional by a constant c to the amount weight at age exceeds the weight at maturation w_m

$$F_a = \begin{cases} c(w_a - w) & \text{if } w_a > w \\ 0 & \text{if } w_a > w \end{cases}. \tag{3.2}$$

where weight is an isometric function of length L_a with the average weight at 100 cm given by w_{100}

$$w_a = w_{100} L_a^3 \times 100^{-6}, \tag{3.3}$$

the factor of 100^{-6} is included to ensure $w_a = w_{100}$ when $L_a = 100$ cm. Length at age L_a follows a von Bertalanffy growth curve with asymptotic length L_∞, and growth rate k

$$L_a = L_\infty (1 - e^{-ka}). \tag{3.4}$$

The total reproduction is determined by the sum of age, genotype, and origin specific fecundities $F_{j,i,a}(g)$ weighted by abundance $N_{j,i,a,t}$

$$N_{j,j,0,t} = \sum_{i \in W,H} \sum_{a=1}^{A_{max}} N_{j,i,a,t} \int_{-\infty}^{\infty} F_{j,i,a}\,(g)\, dg \tag{3.5}$$

the genotype distribution of the yearlings is determined by the distribution of genotypes in the spawning population, the formation of mating pairs, and the transmission of genotypes from parents to offspring. The distribution of genotypes in the spawning population is the sum of genotypes for each age class weighted by their abundance and fecundity

$$\psi_{spawn,j,t}(g) = \frac{1}{Z_{j,t}} \sum_{i \in \{W,H\}} \sum_{a=1}^{A_{max}} N_{j,i,a,t} F_{j,i,a}(g) \tag{3.6}$$

where $Z_{j,t}$ is a normalizing constant. Mating pairs form at random such that the genotypes g_1 and g_2 of the parents of each yearling can be treated as a random sample from the spawning population $\psi_{spawn,j,t}(g)$. The transmission of genotypes follows the infinitesimal model of quantitative genetics (**31**), which assumes genotypes are determined by small contribution from alleles at many loci. Given this model, the genotype g of yearlings with parent genotypes g_1 and g_2 are normally distributed around the mid-parental value with variance V_r. Combining these three assumptions yields an integral equation for the genotype distribution of yearlings

$$\psi_{j,0,t}(g) = \frac{1}{\sqrt{2\pi V_r}} \int_{-\infty}^{\infty} \int_{\infty}^{\infty} e^{\frac{-(g-0.5g_1-0.5g_2\ )^2}{2V_r}} \psi_{spawn,j,t}(g_1)\, \psi_{spawn,j,t}(g_2) dg_1 dg_2 \tag{3.7}$$

(Slatkin 1970).

3.3.2.3. *Domestication selection.* After reproduction, hatchery fish experience stabilizing viability selection. We model viability selection as a single event with survival probability (fitness) determined by the distance from the optimum genotype θ_H and selection strength $s_{v,H}$

$$W_{v,H}(g) = e^{-0.5 s_{v,H} (g-\theta_H)^2}. \tag{3.8}$$

The effect of selection on the genotype distribution of the population is the product of the initial genotype distribution and the survival probability, normalized by $Z_{v,H,t}$

$$\psi'_{H,n,0,t}(g) = \frac{1}{Z_{v,H,t}} \psi_{H,0,t}(g)\, W_{v,H}(g). \tag{3.9}$$

3.3.2.4. *Natural selection.* The wild population experiences two viability selection events: one before and one after density-dependent mortality with survival probabilities determined by stabilizing selection gradients $W_{v_1,W}(g)$ and $W_{v_2,W}(g)$. Following Dédarre and Gandon (2011), selection before density dependence has a smaller effect on the abundance of the population and is called

"soft" selection, while selection after density dependence or "hard" selection has a larger demographic effect. These two selection events have the same optimum genotype θ_W, but different selection strengths $s_{v_1,W}$ and $s_{v_2,W}$ respectively, such that the fitness function is (survival probability) for each is $W_{v_x,w}(g) = e^{-0.5 s_{v_x,w}(g-\theta_w)^2}$. The survival probability integrated over all genotypes determines the proportion of the population that survives selection and therefore the abundance after the first $N'_{W,0,t}$ and second selection events $N'''_{W,0,t}$

$$N'_{W,0,t} = N_{W,0,t} \int_{-\infty}^{\infty} \psi_{W,0,t}(g)\, W_{v_1,W}(g) dg \tag{3.10}$$

$$N'''_{W,0,t} = N''_{W,0,t} \int_{-\infty}^{\infty} \psi''_{W,0,t}(g)\, W_{v_2,W}(g) dg, \tag{3.11}$$

where density dependence (described in the next sub-section below) occurs between these two selection events to determine $N''_{W,0,t}$. The genotype distribution after natural selection is given by the product of the fitness function and genotype distribution, normalized by $Z_{v_x,t}$

$$\psi'_{W,0,t}(g) = \frac{1}{Z_{v_1,t}} \psi_{W,0,t}(g)\, W_{v_1,W}(g) \tag{3.12}$$

$$\psi'''_{W,0,t}(g) = \frac{1}{Z_{v_2,t}} \psi''_{W,0,t}(g)\, W_{v_2,W}(g), \tag{3.13}$$

where again, density dependence determines $\psi''_{W,0,t}$ as described below.

3.3.2.5. *Density dependence.* In white sturgeon, there is evidence that yearlings are most sensitive to density-dependent mortality (**33**). Therefore, we model a density dependence with a single mortality event acting on yearlings between the two selection events. Survival is decreasing function of abundance, consistent with a Beverton Holt stock recruit relationship $f\left(N'_{W,0,t}\right) = \frac{1}{1+\beta N'_{W,0,t}}$, where β determines the strength of density dependence. The abundance of the wild population after density-dependent mortality is

$$N''_{W,0,t} = \frac{S_0 N'_{W,0,t}}{1 + \beta N'_{W,0,t}}. \tag{3.14}$$

Finally, we assume, density-dependent survival is independent of genotype, so the genotype distribution is unchanged $\psi''_{W,0,t}(g) = \psi'_{W,0,t}(g)$.

3.3.2.6. *Immigation.* After density dependence and selection, hatchery-origin individuals emigrate. The number of hatchery-origin fish emigrating from the Nechako River population is subject

to uncertainty and is determined by many factors, including the number of releases, rearing strategies, and variable environmental conditions. To accommodate this uncertainty, we test a range of scenarios determined by a fixed number of immigrants each year when the hatchery is in operation

$$N'''_{H,f,0,t} = \overline{N}_{H,0}. \tag{3.15}$$

We represent the immigration rate as a proportion of natural recruitment at equilibrium $r_H = \overline{N}_{H,0}/N^*_{W,f,1}$ in our results to account for the effect of the wild population size on the rate of gene flow.

3.3.2.7. *Transitions between age classes.* After selection, density dependence, and immigration, yearlings enter the first age class such that

$$N_{i,1,t+1} = N'''_{i,0,t} \tag{3.16}$$

$$\psi_{i,1,t+1}(g) = \psi'''_{i,0,t}(g). \tag{3.17}$$

All age classes one year and older experience density-independent mortality with survival probability S_a before transitioning to the next age class

$$N_{i,a+1,t+1} = S_a N_{i,a,t} \quad \text{for} \quad a > 0. \tag{3.18}$$

Survival between age classes does not depend on genotype; therefore, the genotype distribution of all cohorts (accept yearlings) is unchanged after each time step

$$\psi_{i,a+1,t+1}(g) = \psi_{i,a,t}(g) \quad \text{for} \quad a > 0. \tag{3.19}$$

3.3.3. Genetic parameters. There is a significant body of evidence linking captive breeding to rapid declines in fitness in the wild (**30**). These studies measure the relative reproductive success of captive-raised and wild-origin fish in the wild by comparing the performance of individuals with mixed wild and hatchery ancestry in the wild and hatchery environments, respectively. These studies, however, all require multiple generations of data and, therefore, have only been conducted in short-lived species like salmonids (**7**). Acknowledging the uncertainty associated with using these estimates in a longer-lived species, we calibrate the range of genetic parameters of the model (Table 3.1) to match the fitness declines in early-generation hatchery fish estimated by the meta-analysis

by Christie et al. (2014), as the best available data on the range of plausible values. Christie et al. (2014) observed fitness declines between 30% and 60% in a single generation. These declines are attributed to both genetic (adaptation) and environmental effects (plasticity). The relative contribution of genetic and environmental effects varies by program and species (**6**, **34**). We follow a prior modeling study (**19**) and consider a range of fitness declines caused by genetic effects between 0.0% and 30% per generation, with a baseline value of 15%. We calculate the relative fitness of hatchery-origin fish after one generation of hatchery production by fist solving for the equilibrium genotype distribution of the population at birth under natural selection $\psi_W^*(g)$. We then calculate the genotype distribution of the first generation of hatchery fish $\psi_{g_1}(g)$ by applying the model of reproduction and domestication selection to the equilibrium genotype distribution after natural selection $\psi_W'''^*(g)$. We compute the relative fitness in the wild by comparing the average fitness of each population under natural selection.

$$Q_{g_1} = \frac{\int_{-\infty}^{\infty} e^{-\frac{s_W}{2}(g-\theta_W)^2} \psi_{g_1}(g) dg}{\int_{-\infty}^{\infty} e^{-\frac{s}{2}(g-\theta_W)^2} \psi_W^*(g) dg}, \tag{3.20}$$

where $s_W = s_{f,W} + s_{v_1} + s_{v,2}$ is the total strength of selection summed across the three selection events. We used an optimization algorithm to solve for the strength of the domestication section events ($s_{f,H},\ s_{v,H}$), optimum phenotype in the hatchery (θ_H) and the total strength of natural selection (s_W) which produced the target value of fitness decline. This optimization problem is over-determined because we are tuning four variables to achieve one outcome. Therefore, we add additional constrains by specifying the expected effect of fecundity selection on the average fitness of wild origin fish in the hatchery $\left(S_f = \int \psi_W^*(g) W_{f,H}(g)\, dg\right)$, and the expected survival of the domesticating viability selection event $\left(S_v = \int \psi_{g_1}(g) W_{v,H}(g)\, dg\right)$. We select the parameter combination that met these three constraints with the minimum strength of natural selection s_W and distance between the optimum hatchery genotype θ_H and a target value $\overline{\theta}_H$. We solve the optimization problem using the Nelder-Mead algorithm implemented in the Optimization.jl package in Julia programming language (**35**). The contribution of each natural selection event to the total selection strength s_W determines the effect of selection on demographic outcomes and gene flow. For the base parameter set we fix the value of the soft selection and fecundity selection event at $s_{f,W} = s_{v_1,W} = 0.25 s_w$ with the rest contributed by the hard selection event $s_{v_2,W} = 0.5 s_w$. We consider the effects of alternative parameter choices in the Appendix B.1.

Parameter	Interpretation	Value	range
Q_{g_1}	Relative fitness of fist generation hatchery fish genetic component	0.85	(0.7, 1.0)
Δ	Relative fitness of fist generation hatchery fish environmental component	0.85	(0.7,1.0)
S_f	Effect of fecundity selection on first-generation hatchery fish	0.1	
S_v	Fraction of fish generation hatchery fish that survival viability selection.	0.01	(0.05, 0.005)
V_r	Recombination variance	0.5	
θ_W	Optimal genotype in the wild	0.0	
θ_H	Optimal genotype in hatchery	10.0	
s_W	Total strength of natural selection	0.25	(0.18,0.61)
$s_{f,W}$	Strength of natural fecundity selection	$0.25s_W$	$(0.0s_W, 1.0s_W)$
$s_{v_1,W}$	Strength of soft viability selection	$0.25s_W$	$(0.0s_W, 1.0s_W)$
$s_{v_2,W}$	Strength of hard viability selection	$0.5s_W$	$(0.0s_W, 1.0s_W)$
$s_{v,H}$	Strength of domesticating viability selection	0.10	(0.06, 0.13)
$s_{f,H}$	Strength of domesticating fecundity selection	0.05	

TABLE 3.1. Base value and ranges of the parameters of the genetic model

3.3.4. Demographic parameters. We chose the demographic parameters in our analysis (Table 3.2) to represent the life history characteristics of the mid-Fraser River white sturgeon population. We obtained information on growth, survival, maturation, and fecundity from population monitoring studies and previous demographic modeling exercises (**21**, **24**, **26**). No information was available to estimate yearling survival rates or the strength of density dependence. We tuned these parameters to fix equilibrium recruitment at 100 individuals per year, using the equilibrium conditions given in Appendix B.2, and varied the strength of density dependence. We tuned the strength of density dependence β and the yearling survival S_0 to adjust the strength of density dependence without changing equilibrium recruitment (Appendix B.3). These two changes, in turn, change the basic reproductive number of the population R_0 (i.e. the expected number of offspring produced in an individual's lifetime in the absence of density dependence), with larger values of R_0 implying greater density dependence (higher values of β). To improve the interpretability of our results, we report these values of R_0 to represent the changes in the strength of density dependence.

3.3.5. Simulations. We numerically iterated eqn. (3.1-3.19) to characterize the recipient population's response to immigration by hatchery-origin fish. We developed numerical procedures to approximate the genotype distribution in the Julia programming language (**35**). A detailed

Parameter	Interpretation	Value	range	source
L_∞	Asymptotic or maximum length	300cm		(**24**)
k	Von Bertlanffy growth coefficient	0.034		(**24**)
a_0	Length at age zero	0.0		
w_{100}	Weight at 100cm	9.3Kg		(**24**)
w_m	Weight at maturation	33Kg		(**24**)
c	Size to fecundity conversion factor	1	$w.l.o.g$	
S_a	Survival from age a to age $a+1$	0.96		(**24**)
R_0	Basic reproductive number	3.0	(1.5-9.0)	

TABLE 3.2. Base value and ranges and sources for demographic parameters

Parameter	Interpretation	Value	Range
r_H	Annual proportional hatchery immigration	0.5	(0.0, 1.0)
T	Hatchery duration	100	(0, 250)

TABLE 3.3. Parameter values and ranges used in the simulations

description of these methods is in Appendix B.4. We initiated each simulation with the population at equilibrium without a hatchery, iterated forward in time for $T \in (0, 250]$ years with annual proportional immigration from the hatchery of $r_H \in (0, 1]$, and then simulated another 20 generations without immigration to evaluate the recovery dynamics. For each simulation, we tracked four statistics that summarized the genetic and demographic state of the population (Table 3.3): 1) we characterized changes in abundance by tracking annual recruitment $N_{1,W,t}$, and we described changes in the genotype distribution with 2) the average fitness $\overline{W}_{W,t}$, 3) the mean genotype $\overline{g}_{1,W,t}$, and 4) the genetic variance $V_{1,W,t}$. We extracted the minimum fitness $\overline{W}_{min}$, minimum abundance N_{min} and cumulative immigration $r_H \times T$ from each simulation to characterize the maximum effect of the hatchery program.

3.3.6. Sensitivity analysis. We used a local sensitivity analysis to identify the effects of the genetic and life history parameters on the cumulative fitness effects of straying by hatchery-origin fish. To account for the effects of immigration rate r_H and the duration of the hatchery program in years T we calculated the difference between the equilibrium fitness and the minimum fitness $\Delta\overline{W}(r_H, T) = \overline{W}^*_W - \overline{W}_{min}$ across the range of values of r_H and T. We found this was well approximated by a log-linear function $\Delta\overline{W}(r_H, T) \approx \alpha r_H T^2$, where the parameter α determined the magnitude of the genetic effects for a given combination of r_H and T. We tested the sensitivity of the parameter α to the genetic and life history parameters of the model using a local sensitivity

analysis where the elasticity of α to a parameter θ is defined as $\lambda_\alpha(\theta) = \frac{\theta}{\alpha}\frac{\partial\alpha}{\partial\theta}$. We computed the slope coefficient α of the log-linear relationship using the GLM package in Julia and the partial derivatives with finite differences.

3.3.7. Mitigation measures. We evaluated the effectiveness of two approaches to reducing the fitness effects of the hatchery: 1) reducing the annual proportion of the hatchery-origin fish r_H, 2) reducing the duration of hatchery supplementation T, and 3) removing hatchery-origin individuals after they have immigrated. For hatchery emigration and duration, we characterized the effect of a small reduction in each factor on the minimum fitness ($\overline{W}_{min}$) by calculating the elasticity of $\overline{W}_{min}$ to r_H and T $\left(i.e.\ \frac{r_H}{\overline{W}_{min}}\frac{\partial\overline{W}_{min}}{\partial r_H}, \frac{T}{\overline{W}_{min}}\frac{\partial\overline{W}_{min}}{\partial T}\right)$ across the range of r_H and T values given in Table 3.3. We calculated the derivatives for these elasticity values using the finite difference method. For removal of hatchery-origin strays, we solved for the number of removals required to reduce the fitness loss by 50%. To this end, we defined a function that returns the maximum fitness loss $\mathcal{L}(E,\vec{q}) = \overline{W}^*_w(E,\vec{q}) - \overline{W}_{min}$ as a function of the total removal effort E and the catchability of each age class $\vec{q}$. We applied the bisection algorithm (Roots.jl package) to solve for the level of removal effort E^* required to reduce $\mathcal{L}$ by 50% for a given vector of catchabilities a. We then calculated the total removals by simulating the model with removal effort E^* summing the total removals in each period. The effect of removing individuals was largely determined by their age class. To help interpret these results, we calculated the expected gene flow caused by a single hatchery-origin individual from each age class (Appendix B.5).

3.4. Results

3.4.1. Transient dynamics in fitness effects and demographic outcomes. Our model captures a small decline in recruitment of the wild population (2%) after 100 years of hatchery production (Figure 3.2B) owing to a decline (5%) in fitness (Figure 3.2A). Gene flow from the hatchery shifts the genotype distribution of the wild population away from the local optimum and increased the genetic variance, which together reduces the average fitness (Figure 3.2C,D). We found a substantial delay between when hatchery production occurred and the eventual decline in abundance (Figure 2B). Hatchery stocking can create a demographic boost that leads to increased abundance in the short run, followed by a decline caused by the reduction in fitness. We found that the minimum abundance often occurred over 100 years after hatchery production ended. This

delay between the end of hatchery production and the eventual decline in abundance depended on the strength of density dependence measured by the basic reproductive number. Smaller values of R_0 lead to longer delays between the decline in fitness and the eventual demographic response (Figure B.1) because density dependence is weaker, which strengthens the feedback between the population's reproductive potential and future population sizes. The decline in abundance also depends on the timing of selection, with the largest declines caused by viability selection after density dependent mortality (Figure B.2), the relative contribution of genetic and environmental fitness effects (Figure B.3), and the strength of natural selection (Figure B.4).

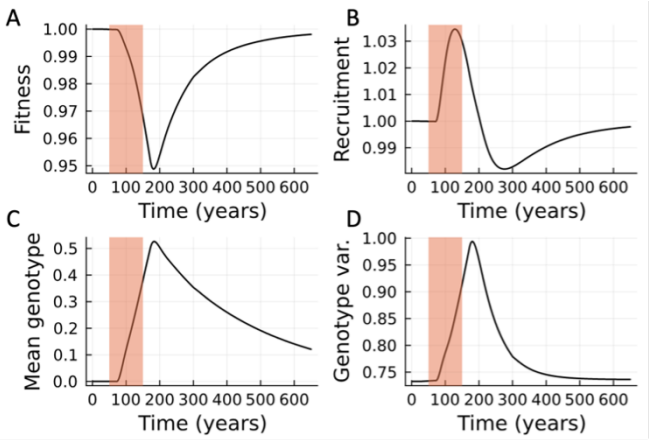

FIGURE 3.2. Time series of the A) fitness, B) recruitment, C) mean genotype and D) genotype variance of the wild population under the default parameter values in Tables 3.1-3.3 The red shading indicates the period of hatchery operation.

3.4.2. Effect of immigration rate and hatchery duration. Minimum fitness and recruitment decreased with increasing hatchery duration (T) and immigration (r_H), and the smallest values of fitness and abundance occurred when both parameters (T and r_H) were large (Figure 3.3A,C). Based on the deterministic simulations, the effect on fitness increased linear with the rate of immigration and with the 3/2 power of the duration of the hatchery program ($T^{3/2}$), while the effect on recruitment increased linearly with the rate of immigration (r_H) and with the square of duration (T^2; Figure 3.3B, D). These exponents with values greater than one imply that increasing hatchery duration has a disproportionally larger effect on outcomes. The convex relationship between hatchery duration and fitness effects occurs because the duration of the hatchery program determines the number of generations in captivity and, therefore, the level of domestication as well

as the cumulative number of immigrants and gene flow to the wild population. In contrast, the rate of immigration only affects the amount of gene flow from the hatchery.

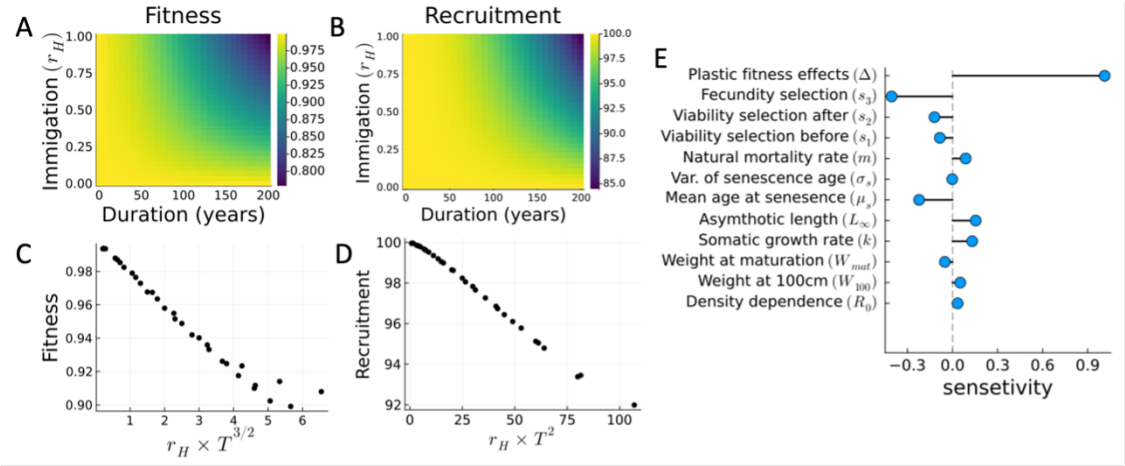

FIGURE 3.3. Time series of the A) fitness, B) recruitment, C) mean genotype and D) genotype variance of the wild population under the default parameter values in Tables 3.1-3.3 The red shading indicates the period of hatchery operation.

3.4.3. Effect of biological parameters on minimum fitness. In the local sensitivity analysis, both the genetic and demographic parameters determine the genetic impacts of the hatchery program (Figure 3.3E). Of the genetic parameters, the plastic effects of the hatchery (Δ) and the strength of fecundity selection (s_3) have the largest effect on the accumulation of fitness effects because these two mechanisms determined the efficiency of gene flow from the hatchery to the wild population. The strength of the viability selection events also reduced the genetic impacts of the hatchery by limiting gene flow. They do so to a lesser degree, however, because viability selection does not act directly on hatchery-origin fish; rather, it affects their progeny. The fitness effects of the immigration by hatchery-origin fish are also determined by the generation time of the population. The demographic parameters with the largest impact were the somatic growth rate (k), asymptotic length (L_∞), age at senescence (μ_s), and the natural mortality rate (m), which interact with each of which determine the generation time. Longer generation times lead to smaller fitness effect. This can be explained in two ways. First, fitness effects accumulate more slowly in long-lived species. Therefore, given a fixed duration of hatchery operation, we expect lower total impacts in longer-lived species. Second, in longer-lived populations, the age structure acts as a

buffer, reducing the genetic impact because each year class contributes less (in proportional terms) to the total reproductive potential of the population.

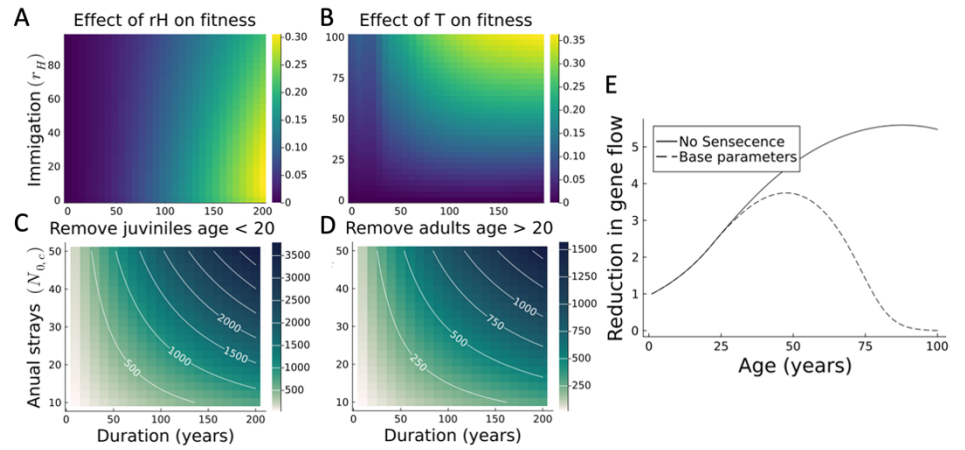

FIGURE 3.4. The effects of mitigation strategies. The marginal effect of reducing the immigration rate from the hatchery A) and the duration of the hatchery B) on the minimum fitness of the wild population. C) The total number of adults (age 20 and over) and D) juveniles (age 19 and under) removed to achieve a reduction in the genetic load by 50%. E) The expected reduction in gene flow from removing an individual as a function of its age for two alternative sets of demographic parameters, one with senescence (dashed line) and one without senescence (solid line).

3.4.4. Effects of alternative mitigation strategies. We found that reducing the duration of hatchery operation and reducing the immigration rate of hatchery-origin fish were similarly effective strategies for mitigating the genetic effects of hatchery production (Figure 3.4A,B), but the effects of a small change in each depended on both the immigration rate and the planned duration of the program. Reducing the immigration rate resulted in the largest increase in fitness when the planned duration of the program was long and when the initial immigration rate was low (Figure 3.4.A). The relationship between the initial immigration rate and the effect of a small change is caused by a saturating relationship between the immigration rate and the minimum fitness of the wild population. As the immigration rate increases, the fitness effects approach a maximum value. Near this point, small reductions in gene flow have a limited effect on the wild population. In contrast, we found that reducing the duration of the hatchery program was most effective when immigration rates were high (Figure 3.4.B). This is likely because when gene flow is high, the domestication of the hatchery population (which increases over time) more directly impacts the wild population. Translocating or removing hatchery-origin individuals can also mitigate the fitness

effect of hatchery releases, but the effectiveness of the strategy depends on the age of the removed individuals (Figure 3.4C, D). Removing older individuals had a larger effect because they had higher expected lifetime reproductive potential (Figure 3.4E). An older mature individual is more likely to reproduce successfully and affect the genetic state of the population. In contrast, a young individual that is removed would have had a higher probability of dying before it reproduced (Figure 3.4E). Removing very old individuals also had a small effect if they had reached the age of senescence.

3.5. Discussion

Our model shows that even temporary hatchery programs can have the potential to create long-lasting effects on the fitness and abundance of wild populations, but the longevity of the species, the strength of domestication selection, and the duration of the program strongly influence the magnitude of those impacts. In our model, gene flow from hatchery-origin individuals that immigrate results in a decline in fitness in the adjacent wild population, an expected outcome given the parameterization based on evidence for domestication and decline in relative fitness of hatchery populations (**30**). Once hatchery operations end, natural selection causes the fitness of the wild population to eventually recover back to the optimum, given the available genetic variance in our model. Over the range of parameter values we considered, gene flow is a stronger force than selection, causing genetic effects to accumulate quickly in the wild population and recover more slowly back to the original fitness, resulting in long-lasting impacts. Despite this potential for long-term impacts, we found that the maximum genetic effects of temporary hatchery programs in our long-lived example species were much smaller than the potential long-run impacts if stocking continued indefinitely (**14**, **16**). Therefore, limiting the duration of captive breeding programs can be an effective strategy for mitigating unintended fitness consequences. Doing so reduces the time required for domestication to occur in the hatchery environment and the total amount of gene flow from the hatchery population to the wild, limiting the overall genetic load. The duration of captive breeding is one of several interacting risk factors identified in the case study of Nechako white sturgeon. The combined effects of the relative fitness of hatchery-origin individuals, the immigration rate, the program duration, and the population's generation time determine the scale of genetic risks. Risks from each of these factors roughly combine in a roughly multiplicative fashion, which means that low levels of risk from only one factor can offset high values of the other three. For

example, careful genetic management and cultivation practices can reduce domestication (**36**, **37**). In these cases, the genetic risk of captive breeding can be small even if it lasts for many generations and straying is high. Conversely, reducing the duration of captive breeding can offset cultivation practices that increase the risk of domestication. For example, currently (in 2023), Nechako River hatchery sturgeon are released at two years of age to increase survival. While this practice might increase domestication selection, the long generation time of the species (white sturgeon mature between ages 20-40 on average) acts to offset these risks. Finally, if logistically feasible, removals from non-target populations can also reduce the rate of gene flow, offsetting high values of immigration, domestication, and duration. We found that the effectiveness of removing individuals increased with their age because the expected reproductive success (and thus gene flow) tends to increase through an individual's life cycle. The population's life history characteristics determined outcomes by changing the time scale of the population dynamics. We found that both the rate of accumulation of genetic load and the rate of recovery are a function of the population's generation time. Willoughby & Christie (2019) found a similar result for the effects of hatcheries on fitness and genetic diversity with an individual-based model by comparing species with different life histories. Measuring the duration of captive breeding in terms of generations (rather than years) can largely account for these effects of life history and differences between species. For example, all else equal, forty years of hatchery production of white sturgeon would have a similar effect to four years for Chinook salmon (*Oncorhynchus tshawytscha*). Re-scaling time in this way allows us to draw a rough equivalence between the dynamics of populations with very different life histories. Interestingly, models of the long-term effects of aquaculture escapees show that long-lived species can be more sensitive to the fitness consequences of gene flow due to greater total accumulation of fitness effects over their lifespan (**38**), further emphasizing the differences between the long-run and transient dynamics of the system.

3.5.1. Implications for adaptive management. The time scale of population dynamics relative to the time scale of management interventions can determine the effectiveness of such interventions. In the case of the Nechako River white sturgeon system, a mismatch in these time scales might provide an opportunity. Because fitness effects accumulate slowly, captive breeding can maintain recruitment in the Nechako River population for a long period with relatively low risks to neighboring populations, thus providing time to identify and address the root causes of the

recruitment failure. In shorter-lived species, restoration efforts might take the same amount of time as they would for longer-lived species due to similar levels of scientific uncertainty and bureaucratic delays. The genetic effect of captive breeding, however, would accumulate more quickly. Longevity, therefore, can have a strong effect on the balance between conservation benefits from hatchery production and genetic risks. The time scale of population dynamics can also determine the ability of managers to learn and adapt. In the first years of operation, the Nechako River hatchery released large numbers of yearling sturgeon, observed high natural mortality, and detected emigration to the mid-Fraser River. In response, the program shifted to releasing fewer larger individuals, causing an apparent decline in both mortality and emigration. In a short-lived population, these initial year classes could have caused larger decreases in fitness in the mid-Fraser River if they interbred before managers could have observed and responded to the event. In the Nechako River Sturgeon system, however, the impact of these cohorts is likely to be relatively small because each year class contributes less to the total reproductive potential of the population. In this case, each year provides an opportunity to learn, and because the population dynamics are slow, adaptive management approaches (*sensu.* (**39**)) can be employed quickly relative to the rate at which impacts accumulate.

3.5.2. Limitations and Assumptions. As with all modeling studies, ours involves simplifying assumptions, and the interpretation of our results depends on this context. The quantitative results of our model are sensitive to the value chosen for the fitness declines of hatchery-origin fish in the natural environment. There is a substantial body of evidence suggesting that differences in relative fitness can arise from both domestication selection and plastic environmental effects (**6**,**7**). The cumulative fitness declines and the relative contribution of genetic and environmental effects, however, varies between systems (**7**), and measurements of these quantities are not available in our study system. This limits the quantitative precision of our model estimates. Our qualitative findings about the relative effectiveness of the alternative mitigation measures and the effects of life history traits, however, are robust to this source of uncertainty. Furthermore, the scenarios with high estimated fitness effects likely represent an upper bound on the program's potential impact, given that much of the uncertainty is associated with potential mitigating factors such as the timing of natural selection and the relative contribution of environmental effects.

In addition to parametric uncertainty, our model uses a simplified representation of the population's genetics. We assume the effects of domestication can be represented by a single quantitative trait. This model captures the additive fitness effects of domestication selection, but it also assumes there is sufficient genetic variation for the population to recover back to the wild phenotypic optimum once hatchery production ends. In a small population, however, unfavorable alleles might reach fixation through genetic drift, permanently limiting the population's ability to recover (**40**, **41**). Furthermore, if multiple traits are under selection, differences in selection strengths and heritabilities between these traits can result in quantitative changes in the accumulation of genetic load accumulation and recovery. Still, our single-trait model likely captures the essential qualitative features of these processes (**42**, **43**).

Our model also makes several simplifying assumptions about the population's ecology. We assume that the primary ecological interactions between hatchery and natural origin individuals occur before recruitment. This is consistent with a meta-analysis that shows density-dependent mortality declines in fish populations with age (**44**). It does not, however, account for other interactions which might operate through the life cycle. For example, competition between juveniles and adults could reduce growth, maturation rates, and fecundity. The presence of hatchery-origin fish in a population can also increase the risk of disease and attract predators (**10**). Incorporating these processes was beyond the scope of our study, but they might play an important role in the demographic effects of hatchery production.

Finally, our model presents a highly simplified representation of the spatial and movement ecology of our study organism. We represent the hatchery and wild populations in our model as two distinct populations separated by a clearly defined boundary; however, these boundaries might be less clear in many real-world systems with multiple populations characterized by patterns such as isolation by distance. Furthermore, we assume that all movement occurs before maturation and that fish do not exhibit homing behaviors, returning to their river of origin to reproduce. It is currently unknown if the hatchery-origin fish that have emigrated from the Nechako River will remain in the mid-Fraser River throughout their life cycle or return to the Nechako to reproduce. If some hatchery fish do return to the Nechako River, gene flow impacts will be reduced; however, ecological impacts from hatchery fish using the mid-Fraser River habitat would remain. More detailed information about

habitat uses and spawning behavior would help reduce uncertainty about the extent of the genetic and ecological effects of the hatchery program.

3.6. Conclusions

In our modeled case study, the long-lived life history of white sturgeon caused the genetic effects of hatchery production to accumulate slowly over time. Given that captive breeding is expected to be a temporary measure in this system, the slow time scale of the population's dynamics is a dominant factor determining the level of genetic risks. This highlights the general importance of time scales and the management of ecological systems. Models that only account for a population's biology often explicitly ignore the time scale of the system's dynamics by studying equilibrium behaviors and rescaling (non-dimensionalizing) the units of time. Although these can be useful methods in some cases, in an applied setting, the time scale of the ecological system's dynamics can influence the options available to managers. In our case study, the relatively slow time scale of the population's dynamics provides a long window of time for managers to learn and address the root causes of the population's decline, with a relatively limited risk of unintended genetic consequences when compared to shorter-lived species. In effect, the time scale of the biological system becomes important in an applied context because of its interaction with the time scale of the management system.

Bibliography

1. K. A. Naish, *et al.*, *Advances in Marine Biology* (Elsevier, 2007), vol. 53, pp. 61–194. DOI: 10.1016/S0065-2881(07)53002-6.
2. S. Kitada, *Reviews in Aquaculture* p. raq.12418 (2020).
3. K. A. Lamothe, *et al.*, *Environmental Reviews* **27**, 575 (2019). ISBN: 1181-8700 publisher: NRC Research Press.
4. A. L. George, *et al.*, *Fisheries* **34**, 529 (2009).
5. L. Laikre, M. K. Schwartz, R. S. Waples, N. Ryman, *Trends in Ecology Evolution* **25**, 520 (2010).
6. H. Araki, B. Cooper, M. S. Blouin, *Science* **318**, 100 (2007).
7. M. R. Christie, M. L. Marine, R. A. French, M. S. Blouin, *Proceedings of the National Academy of Sciences* **109**, 238 (2012).
8. N. Ryman, L. Laikre, *Conservation Biology* **5**, 325 (1991).
9. R. S. Waples, K. Hindar, S. Karlsson, J. J. Hard, *Current Zoology* **62**, 617 (2016).
10. K. Kostow, *Reviews in Fish Biology and Fisheries* **19**, 9 (2009).
11. M. M. McClure, *et al.*, *Evolutionary Applications* **1**, 356 (2008).
12. S. M. Carlson, W. H. Satterthwaite, *Canadian Journal of Fisheries and Aquatic Sciences* **68**, 1579 (2011).
13. A. G. Dedrick, M. L. Baskett, *The American Naturalist* **192**, E62 (2018).
14. M. J. Ford, *Conservation Biology* **16**, 815 (2002).
15. S. Edmands, C. C. Timmerman, *Conservation Biology* **17**, 883 (2003).
16. M. L. Baskett, R. S. Waples, *Conservation Biology* **27**, 83 (2013).
17. N. F. R. Snyder, *et al.*, *Conservation Biology* **10**, 338 (1996). Publisher: [Wiley, Society for Conservation Biology].
18. S. E. McMurray, K. J. Roe, *Freshwater Mollusk Biology and Conservation* **20**, 1 (2017). ISBN: 2472-2944 publisher: BioOne.
19. J. R. Willoughby, M. R. Christie, *Conservation Biology* **33**, 377 (2019).
20. COSEWIC, Cosewic assessment and status report on the white sturgeon acipenser transmontanus in canada. committee on the status of endangered wildlife in canada (2012).
21. L. RLL Environmental Services, Fraser river white sturgeon monitoring program: comprehensive report (1995 to 1999). (2000).
22. A. D. Schreier, B. Mahardja, B. May, *Canadian Journal of Fisheries and Aquatic Sciences* **69**, 1968 (2012).
23. Fisheries, O. Canada, Recovery strategy for white sturgeon (acipenser transmontanus) in canada [final] (2014).

24. E. Smyth, J. Abell, B. Toth, T. Hatfield, Recovery potential assessment for white sturgeon (acipenser transmontanus) upper fraser designatable unit (2016).
25. J. Korman, C. Walters, Nechako river white sturgeon recovery planning: Summary of stock assessment and oct. 2-3, 2000 workshop (2001).
26. C. C. Wood, D. Sneep, S. McAdam, J. Korman, T. Hatfield, Recovery potential assessment for white sturgeon populaitons listed under the species at risk act (2007).
27. P. by Susan M. Pollard., Breeding Plan for Nechako White Sturgeon. (2005).
28. T. C. Nelson, D. Robichaud, T. Mochizuki, T. Thibault, Fraser river region 3 white sturgeon population monitoring, 2020 (2021).
29. J. Tufto, *The American Naturalist* **158**, 331 (2001).
30. M. R. Christie, M. J. Ford, M. S. Blouin, *Evolutionary Applications* **7**, 883 (2014).
31. N. Barton, A. Etheridge, A. Véber, *Theoretical Population Biology* **118**, 50 (2017).
32. F. Débarre, S. Gandon, *The American Naturalist* **177**, E84 (2011).
33. C. Justice, *et al.*, *Canadian Journal of Fisheries and Aquatic Sciences* **66**, 802 (2009).
34. K. S. Williamson, A. R. Murdoch, T. N. Pearsons, E. J. Ward, M. J. Ford, *Canadian Journal of Fisheries and Aquatic Sciences* **67**, 1840 (2010).
35. J. Bezanson, A. Edelman, S. Karpinski, V. B. Shah, *SIAM Review* **59**, 65 (2017).
36. M. A. Hess, C. D. Rabe, J. L. Vogel, J. J. Stephenson p. 15 (2012).
37. I. Janowitz-Koch, *et al.*, *Evolutionary Applications* **12**, 456 (2019).
38. L. Yang, R. S. Waples, M. L. Baskett, *Theoretical Population Biology* **129**, 93 (2019).
39. C. J. Walters, C. S. Holling, *Ecology* **71**, 2060 (1990).
40. R. Lande, *Researches on Population Ecology* **40**, 259 (1998).
41. M. Castellani, *et al.*, *Evolutionary Applications* **11**, 1010 (2018).
42. J. Tufto, *Evolution* **64**, 180 (2010).
43. J. Tufto, *Evolution* **71**, 2262 (2017).
44. K. Lorenzen, E. V. Camp, *Fisheries Research* **217**, 5 (2019).

CHAPTER 4

Investing in Information for Fisheries Management

4.1. Abstract

Successful conservation and natural resource management often require significant investments in monitoring programs to allow management actions to respond to changes in the ecosystem. However, financial and scientific resources for these programs are often limited and must be allocated between competing priorities. This problem is particularly acute for fisheries management, where harvest limits are ideally updated to track changes in population abundance, but in many cases, the number of stocks in need of assessment exceeds the available capacity. I developed a bioeconomic model to identify the optimal strategy for investing in monitoring information used to set catch limits for a fishery. Using this model, I tested the sensitivity of the economic value of monitoring information, such as fisheries independent surveys and model-based stock assessments, to the manager's objectives and the population's biology. Based on these findings, I identified a set of factors that can be used to prioritize scarce monitoring resources. The relative abundance and productivity of the stock increased the economic value of monitoring information, while stocks with more variable dynamics were assessed more frequently under the optimal strategy. Furthermore, monitoring was more valuable for maximizing fishery profits than for achieving conservation goals. These results follow from the effect of monitoring on the dynamics of the management system. Monitoring allows harvest limits to be updated in response to changes in the stock's status, creating stabilizing feedback that maintains the stock in a productive, economically valuable state. Increasing the frequency of monitoring tightens these feedback loops, which disproportionately increase the value of fisheries targeting stocks with high and variable growth rates whose abundance varies more rapidly over time.

4.2. Introduction

Monitoring programs can be critical components of conservation and natural resource management strategies. Monitoring forms the basis of adaptive management strategies that use monitoring to estimate the effectiveness of past interventions (**1**), and monitoring support state-dependent policies that apply interventions based on current biological and environmental conditions (**2**). More generally, accurate and timely information about the state of the ecosystem allows people to change their behaviors to respond to changes in the underlying ecosystem, establishing a stabilizing feedback mechanism that can increase the resilience of the coupled human-natural system (**3**).

Despite these potential sources of value, the cost-effectiveness of many environmental monitoring programs has been called into question (**2**,**4**). Monitoring consumes large fractions of many conservation and resource management agency budgets (**2**), and these funds might have a larger impact if spent on active interventions like restoration projects or other forms of information collection like basic research. Furthermore, in some contexts, resources for monitoring are scarce, creating a need to prioritize scientific resources for data collection and analysis between systems (**5**).

Fisheries management exemplifies both the potential value of monitoring programs and the need for prioritization. Many successful models of fisheries management require accurate estimates of the abundance of harvested populations to set catch limits, season lengths, and other restrictions on the fishery (**6**). Annual catch limits exemplify the value of state-dependent management. Policies that use a fixed catch limit (e.g., the maximum sustainable yield) are inherently unstable (**7**), but if these limits are adjusted in response to changes in the abundance of the stock, they can be robust to uncertainty and environmental change (**8**). However, given the remote location of many fisheries and other logistical challenges, fisheries monitoring programs that assess abundance levels are a significant expense. In the United States, the National Marine Fisheries Service (NMFS) spent $175 million on fisheries data collection, surveys, and stock assessments in 2021, about 30% of its total fisheries science and management budget ($620 million, NOAA Blue Book 2021). Globally, management agencies often have many more stocks within their portfolios than they can regularly monitor and assess (**5**,**9**), with resources being especially limited in developing countries (**10**). This combination of factors has created a need for fisheries management agencies to develop strategies for prioritizing scientific resources between fisheries based on biological characteristics, economic value, and perceived management needs (**5**).

A critical component of these prioritization decisions is identifying biological attributes and management contexts that determine the value of monitoring information for a given fishery. These factors can provide a science-based approach for allocating scientific resources between competing priorities and for evaluating the efficiency of fisheries monitoring programs. For example, Methot (2015) developed a framework that assigns a numerical score to stocks managed by NMFS based on biological characteristics, conservation goals, and economic value. These numerical scores were then used to identify the target number of years between assessments for each stock and to set a list of annual priorities. Although the choice of factors included in this framework is well motivated by

verbal arguments, more recent advances in bioeconomic modeling techniques provide a set of tools to produce rigorous quantitative estimates of the value of monitoring information and the optimal investment strategy (**11**, **12**, **13**).

Here, I build on these advances in bioeconomic modeling by developing a model of fisheries monitoring programs that characterizes the optimal strategy for investing in monitoring information and estimates its economic value for fisheries management. However, the modeling alone cannot quantify the effect of the biological and economic attributes of the fishery on the value of monitoring. To estimate these values, I combined the bioeconomic model with global sensitivity analysis (**14**), which produces a novel measure of the effect of biological and management context on the value of information. The resulting estimates can help identify when fisheries monitoring programs are cost-effective and prioritize monitoring resources between fisheries when they are scarce. Furthermore, the sensitivity analysis method is highly generalizable, allowing this novel application of the technique to be applied to other systems where monitoring resources need to be prioritized.

4.3. Methods

4.3.1. Model overview. I develop a bioeconomic model of a fishery that tracks the biomass of the stock B_t and the harvest H_t at each annual time step t. A manager sets a catch limit $\overline{H}_t$ each period based on an estimate of the stock's abundance $\hat{B}_t$ and the corresponding level of uncertainty CV_t, derived from noisy observations of the stock y_t. The manager can choose to increase the accuracy of the observations y_t each period by investing in monitoring i_t. The abundance estimate is updated each period in a two-step process. The first step updates the estimate to account for harvest and population growth. The second step updated the estimates using Baye's law to account for the monitoring observation y_t. The manager sets the catch limits $\overline{H}_t$ according to a predefined harvest control rule and chooses when to monitor the stock to maximize an objective function, which includes profit from harvesting the stock, non-consumptive value of the stock, and the cost of monitoring the population.

The model is constructed from three sub-models (fig. 4.1) mirroring the management strategy evaluation (MSE) framework (**15**): 1) an operation model that describes the dynamics of the stock, the harvest H_t, and the monitoring observations y_t, 2) an estimation model, that converts the observations y_t into the abundance estimate, and 3) a management strategy that defines how

the abundance estimate is used to set the target harvest level $\overline{H}_t$ and monitoring decision i_t. This allows me to analyze the dynamics of the stock's abundance, the information available to the manager and their choice of actions as a coupled system. An example of how the dynamics play out is shown in the time series in figure 4.1.B. The manager's abundance estimates track the true abundance of the stock, with uncertainty growing between periods with monitoring. However, in contrast with the MSE framework, which compares the performance of predefined policies, I use dynamic optimization techniques solve for the optimal monitoring decisions.

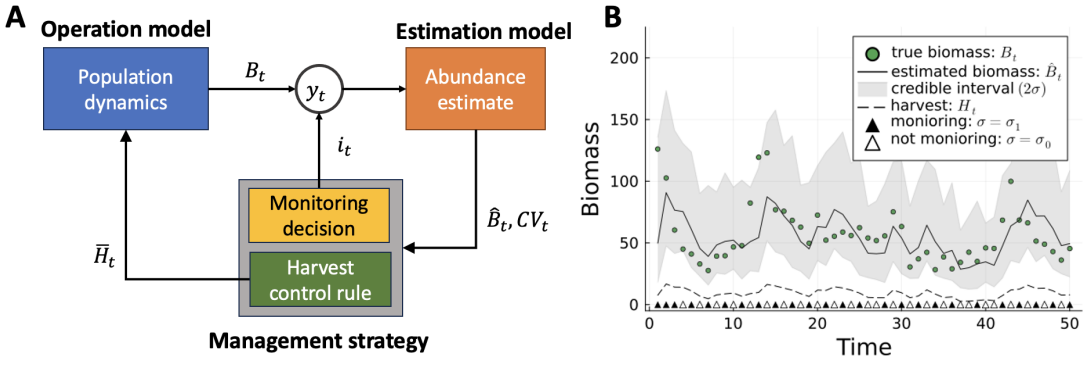

FIGURE 4.1. A) Schematic model diagram illustrating the causal links between the state of the stock, the managers beliefs about the stocks abundance and the resulting harvest and monitoring decisions. B) an example time series of model output, tracking biomass, harvest, monitoring, and the estimated abundance of the stock.

4.3.2. Operation model. The operation model describes how management decisions map to biological and economic outcomes. The model has two management decisions: the target harvest level $\overline{H}_t$ and the monitoring decision i_t. These choices determine the harvest in the current period H_t and the biomass in the following period B_{t+1} through the stock growth equations and the accuracy of the monitoring observation y_t through the observation model.

4.3.2.1. *Stock growth.* I assume the current biomass of the stock B_t depends on the biomass in the previous period B_{t-1}, the harvest limit $\overline{H}_{t-1}$, and variable environmental conditions ν_{t-1}. For each period, the biomass is updated in two steps. First, the harvested individuals are removed from the population. This number generally equals the harvest limit $\overline{H}_{t-1}$, but because the biomass B_{t-1} is uncertain, the harvest limit can exceed the stock biomass. Fisheries rarely, if ever, cause harvest populations to go extinct (**16**), so I cap the harvest level so that it does not deplete the

stock below a minimum biomass B_{min}

$$B'_{t-1} = \begin{cases} B_{t-1} - \overline{H}_{t-1} & \text{if } B_{t-1} - \overline{H}_{t-1} > B_{min} \\ B_{min} & \text{if } B_{t-1} - \overline{H}_{t-1} \leq B_{min} \end{cases} \tag{4.1}$$

I set B_{min} to 5% of the biomass that produces the maximum sustainable yield in the main text and test the sensitivity of our results to this assumption in appendix C.1.1.

Second, the population grows. I assume the growth is a density-dependent function of the biomass B'_{t-1} and environmental conditions ν_{t-1}. For simplicity, I use a two-parameter model consistent with logistic growth, where r describes the growth rate at low biomass and b is the strength of density dependence

$$B_t = \frac{rB'_{t-1}e^{\nu_{t-1}-0.5\tau^2}}{1 + bB'_{t-1}}. \tag{4.2}$$

I assume that the environmental term (ν_t) is stochastic and defined by a sequence of independent and identically distributed normal random variables with mean zero and variance τ^2. These equations define the probability of transitioning $T(B_{t-1}, B_t|\overline{H}_{t-1})$ from biomass B_{t-1} to B_t given the harvest limit $\overline{H}_{t-1}$.

4.3.2.2. *Observation model.* At the beginning of each period, the manager makes an observation y_t of the stock's biomass. Following the previous literature on optimal environmental monitoring, I assume the observations have log-normally distributed measurement errors (**12**, **13**). This choice of distribution ensures the observations will also take a positive value and causes the errors to scale proportionally with the stock size. The variance parameter of the measurement errors $\sigma_{i_{t-1}}$ is determined by the monitoring decision $i_{t-1} \in \{0, 1\}$ made in the previous period. If the manager chooses to invest in monitoring ($i_{t-1} = 1$), the measurement error is σ_1. However, these investments might only represent part of the information used to manage the fishery. For example, $i_t = 1$ can represent investing in the fishery-independent surveys. In this case, the manager might also estimate the stock's abundance using fishery-dependent data. The quality of these alternative information sources is represented by the measurement errors σ_0. Given these assumptions the likelihood of observing y_t given the biomass B_t and monitoring decision i_{t-1} is given by a log-normal probability density function

$$g(y_t, B_t, i_{t-1}) = \frac{1}{y_t \sigma_{i_t} \sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{log(y_t)-log(B_t)}{\sigma_{i_{t-1}}}\right)^2}. \tag{4.3}$$

4.3.3. Estimation model. The biomass estimates are formed using Bayesian inference. The information available to the manager about the stock's abundance is represented by a probability density function $f(B_t)$. Each period, this distribution is updated in a two-step process. First, the initial belief state $f(B_t)$ is updated to account for the new observation y_t using Bayes' law

$$f'(B_t) = \frac{g(y_t, B_t, i_{t-1}) f(B_t)}{\int_0^\infty g(y_t, u, i_{t-1}) f(u) du}. \tag{4.4}$$

Second, the belief state is updated to account for the effects of harvest and population growth. The belief states $f(B_{t+1})$ after harvest, and growth is found by integrating the state transitions probabilities $T(B_{t+1}, B_t|\overline{H}_t)$ over the current belief state $f'(B_t)$

$$f(B_{t+1}) = \int_0^\infty f'(u) T(B_{t+1}, u|\overline{H}_t) du. \tag{4.5}$$

For the optimization procedure, I approximate the manager's belief state with a normal probability density function ϕ with mean $\hat{B}_t$ and coefficient of variation CV_t. I track changes in the belief state by calculating the effects of equations 4.4 and 4.5 on the mean $\hat{B}_t$ and coefficient of variation CV_t as described in Appendix C.2.1. This method is equivalent to the density projection approach used by (**17**).

4.3.4. Management strategy. The management strategy specifies two decisions in each period: 1) the harvest limit $\overline{H}_t$ and 2) monitoring i_t. Both quantities are determined by policy functions that map from the manager's belief state ($\hat{B}_t$, CV_t) to a decision. The catch limit $\overline{H}_t$ is set by a predefined function shown in figure 4.2. This function is based on HCRs commonly used for marine fisheries management in the United States (**18**).

In contrast, I do not set a predefined function for the monitoring decision i_t. Instead, I define a management objective and apply an optimization procedure to solve the policy function that maximizes it. The optimal solution specifies the optimal sequence of actions ($\{i\}_t$) over time, given the uncertainty over the population's biomass. This type of decision problem is called a Partially

Observable Markov Decision Process (POMDP); see (**19**) for a detailed discussion of the application of POMDPs to environmental monitoring.

4.3.4.1. *Harvest control rule.* The harvest control rule is sets the harvest limit $\overline{H}_t$ for each period based on the expected biomass $\hat{B}_t$. The control rule first determines the annual target exploitation ratio $\overline{F}_t$ based on the expected biomass based on the hockey stock rule commonly used in US marine fisheries management (**18**). Above a target biomass level B_{target}, the target exploitation ratio is set at a constant value F_{target}. Below B_{target}, the exploitation ratio declines linearly down to a threshold value B_{limit} below which the target exploitation ratio is set to a minimum value F_{dm}. This is defined as a piecewise function

$$\overline{F}_t(\hat{B}_t) = \begin{cases} F_{target} & \text{if } \hat{B}_t > B_{target} \\ Z_t F_{target} + (1 - Z_t) F_{dm} & \text{if } B_{limit} < \hat{B}_t \leq B_{target} \\ F_{dm} \text{ if } \hat{B}_t \leq B_{limit} \end{cases} \tag{4.6}$$

$$Z_t \equiv \frac{\hat{B}_t - B_{limit}}{B_{target} - B_{limit}}. \tag{4.7}$$

The two threshold values B_{target} and B_{limit} correspond to target and limit reference points (**18**). Figure 4.2 illustrates the effect of these reference points on the control rule. Throughout, we fix the biomass target at the biomass at maximum sustainable yield $B_{target} = B_{MSY}$, the target fishing mortality at the maximum sustainable yield $F_{target} = F_{MSY}$, the limit reference point B_{limit} at 25% B_{MSY} and the *de minimus* fishing mortality F_{dm} at 2%. I also consider a case with a constant fishing mortality target, which corresponds to setting the target and limit reference points to zero ($B_{target} = B_{limit} = 0.0$).

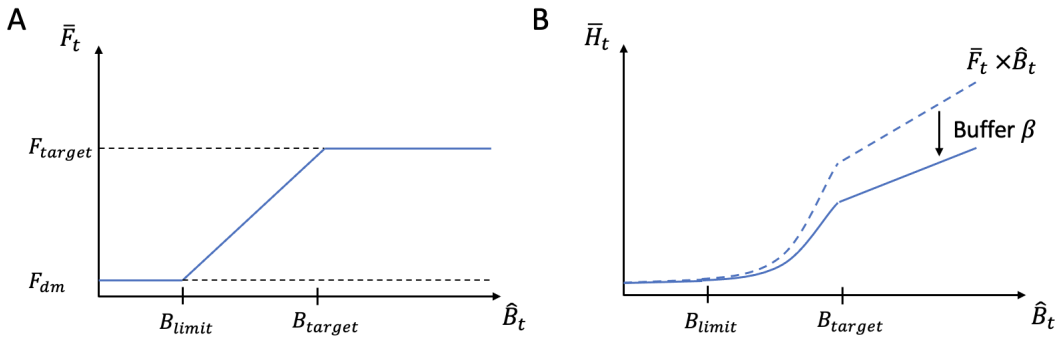

FIGURE 4.2. A) The harvest control rule, relating the target fishing mortality rate $\overline{F}_t$ to the estimated biomass $\hat{B}_t$ based on the reference points B_{limit} and B_{target} and mortality rate targets F_{dm} and F_{target}. B) The harvest limit $\overline{H}_t$ is determined by the target exploitation ratio times the estimated biomass ($\overline{F}_t \times \hat{B}_t$) minus a buffer β.

The harvest limit is set by solving for the harvest that will yield the annual target exploitation ration $\overline{F}_t$ given the expected biomass $\hat{B}_t$, minus a buffer $\beta \in (0, 1)$ to account for uncertainty

$$\overline{H}_t = (1 - \beta)\hat{B}_t\overline{F}_t. \tag{4.8}$$

In the base specification, we assume β is a fixed value that does not depend on the level of uncertainty CV_t. However, some fisheries are managed with harvest control rules that reduce harvest limits when the stock abundance is more uncertain. A common approach is the $p^* - \sigma$ rule (**20**). We explore the effects of this specification in appendix C.1.2.

4.3.4.2. *Manager's problem.* I assume the manager chooses whether or not to monitor the stock each period to maximize profits from the fishery $\pi_H(H_t, B_t)$ and non-consumptive values $\pi_N(B_t)$ minus the cost of monitoring $c \times i_t$. These three components are combined into an objective function that depends on the harvest H_t and biomass B_t, which are given, as well as the choice to monitor i_t, which is chosen optimally.

The profits from the fishery π_H are determined by revenues from harvest, stock-dependent costs, and a quadratic term associated with processing and handling the catch

$$\pi_H(H_t, B_t) = \frac{H_t}{MSY} - \frac{c_1 H_t}{F_{MSY} B_t} - \frac{c_2 H_t^2}{MSY^2} \tag{4.9}$$

The revenues H_t/MSY equal the catch as a fraction of the maximum sustainable yield, and the cost terms are constructed to simplify to c_1 and c_2 when the stock is harvested at the maximum

sustainable yield ($H_t = MSY$ and $F_t = F_{MSY}$). This parameterization facilitates comparison between scenarios with different stock sizes and productivity. It also ensures that the expected profit from harvest set by the HCR is always positive.

I assume the non-consumptive values are a concave function of the stock's abundance to capture the conservation benefits from leaving biomass in the ecosystem and risk aversion

$$\pi_N(B_t) = \frac{\rho B_t}{B_t + (1-\rho)2B_{MSY}}. \tag{4.10}$$

The strength of risk aversion is determined by a parameter ρ that controls the curvature of the function. The non-consumptive values are scaled to equal one when the population is at the equilibrium abundance in the absence of fishing.

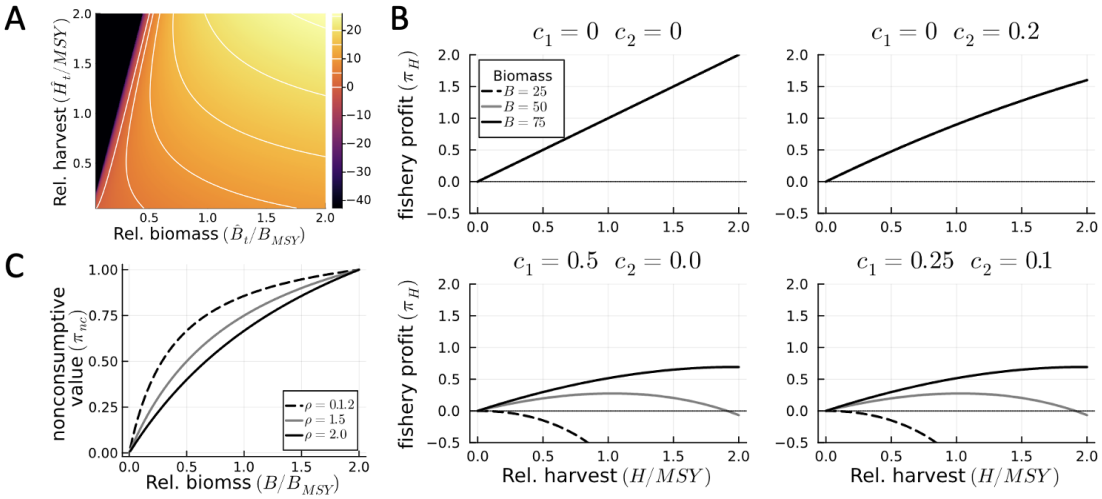

FIGURE 4.3. The effect of the cost and risk aversion parameters on the objective function. A) The value of the objective function in the base parameter set as a function of harvest H_t and biomass B_t. B) The effect of the cost parameters c_1 and c_2 on the fisheries profits as a function of the biomass B_t and harvest relative to the maximum sustainable yield. C) The effect of risk aversion parameter ρ on the non-consumptive values as a function of the biomass B_t relative to the biomass that produces the maximum sustainable yield B_{MSY}.

The manager's problem maximizes the weighted sum of the fisheries profits, non-consumptive values, and monitoring costs c, discounted into the future at rate r, subject to the harvest control

rule, harvest constraints, and the stock dynamics

$$\max_{i_t \in \{0,1\}} \left\{ \sum_{t=0}^{\infty} \left(\frac{1}{1+r} \right)^t \left(\omega_H \pi_H(H_t, B_t) + \omega_N \pi_N(B_t) - c i_t \right) \right\} \tag{4.11}$$

$$s.t. \text{ eqn. } 4.1 - 4.10,$$

where ω_H is the weight given to fisheries profits and $omega_N$ is the weight given to nonconsumptive value.

The solution to the manager's problem is characterized by two components: a policy function $P(\hat{B}_t, CV_t)$ that maps from the manager's belief state to the optimal monitoring decision i_t and a value function $V(\hat{B}, CV_t)$ that maps from the manager's belief state to the expected net present value of the fishery under optimal monitoring. I solved for approximations of these functions using a dynamic programming algorithm described in appendix C.2, and tested their performance using Montecarlo integration in appendix C.3.

4.3.5. Value of information sensitivity analysis. I used the bioeconomic model to compute the effects of the biological and management context on the value of monitoring information. This was done in two stages. First, I defined three metrics that characterize the value of information and the optimal policy: The total value of the monitoring program (VoI^T), the marginal benefit of monitoring net of monitoring costs (VoI^M), and the frequency of monitoring ($\omega_{i=1}$). I then applied a global sensitivity analysis (GSA) to characterize the effects of changes to the model parameters in the bioeconomic model on the value of the monitoring program.

4.3.5.1. *Total value of monitoring.* I calculated the total economic value of the monitoring to understand the importance of monitoring for the performance of the fishery management strategy. I computed this value by comparing the expected net present value of the fishery when monitoring is chosen optimally $V(\hat{B}_t, CV_t)$ to the expected net present value of the system in the absence of monitoring $V_{i_t=0}(\hat{B}_t, CV_t)$. I computed $V_{i_t=0}(\hat{B}_t, CV_t)$ by solving the manager's problem (eqn.4.11) with the value of i_t fixed at zero. The total value of monitoring is given by the difference between these two value functions,

$$VoI^T = V(\hat{B}_t, CV_t) - V_{i_t=0}(\hat{B}_t, CV_t). \tag{4.12}$$

Parameter	Interpretation	base value	range
$\hat{B}_t$	Estimated biomass	N.A.	$0.05B_{MSY} - B_{MSY}$
CV_t	Uncertainty in $\hat{B}_t$	N.A.	$0.0 - 1.0$
MSY	Maximum sustainable yield	10.0	$w.l.o.g.$
F_{MSY}	Exploitation ratio that produces the maximum sustainable yield	0.2	$0.05 - 0.3$
B_{MSY}	The biomass that produces the maximum sustainable yield	50	$33.3 - 200$[1]
r	The population growth rate when rare	N.A. [1]	N.A.[1]
b	The effect of density dependence on population growth	N.A. [1]	N.A.[1]
τ^2	Variance of stochastic environmental conditions	0.05	$0.025 - 0.1$
β	Uncertianty buffer	0.4	$0.0 - 0.3$
B_{min}	Minimum biomass	$0.05B_{MSY}$	N.A.[2]
F_{target}	Maximum target exploitation ratio	F_{MSY}	N.A. [2]
F_{dm}	Minimum target exploitation ratio	0.02	N.A. [2]
B_{target}	Target reference point rule	B_{MSY}	N.A.[2]
B_{limit}	Limit reference point	B_{MSY}	N.A. [2]
$\sigma^2_{i_t=1}$	Variance of observation error when actively monitoring	0.2	$0.1 - 0.5$
$\sigma^2_{i_t=0}$	Variance of passive observations	2.0	$1.0 - 3.0$
c	cost of active monitoring	1.0	$w.l.o.g$
c_1	Stock dependent costs param.	0.1	$0.0 - 0.5$
c_2	nonlinear costs param.	0.05	$0.0 - 0.1$
ρ	Nonconsumptive values nonlinear param.	1.5	$1.2 - 2.0$
ω_H	Harvest values weight	10.0	$1.0 - 15.0$
ω_N	Nonconsumptive values weight	10.0	$1.0 - 15.0$
δ	Discount rate	0.05	$0.01 - 0.1$

TABLE 4.1. Model parameters and ranges used in the global sensitivity analysis.
[1] value is determined by choice of MSY and F_{MSY}
[2] not included in global sensitivity analysis

4.3.5.2. *Marginal benefit of monitoring.* In addition to calculating the total value of monitoring I computed the marginal benefit of choosing to monitor once without cost when otherwise following the optimal policy VoI^M. The goal of calculating this quantity is to understand what factors make monitoring most valuable at a given time, guiding short-term prioritization decisions, like the annual stock assessment priorities described in (**5**). I computed this value by comparing the expected present value of the system assuming monitoring was chosen $i_t = 1$ verses assuming it was not $i_t = 0$. I compute this value using the Q function (appendix C.3), which is the net present value of the fishery given the current belief state $(\hat{B}_t, CV_t)$ and choice of action i_t. Given the Q

function, the expected benefit of a single monitoring observation is found by taking the difference between the net present value when choosing to monitor and not choosing to monitor in the current period, plus the monitoring costs

$$VoI^M(\hat{B}_t, CV_t) = Q(\hat{B}_t, CV_t, 1) - Q(\hat{B}_t, CV_t, 0) + c. \tag{4.13}$$

In addition to the raw value VoI^M used in the sensitivity analysis, I also calculated the annuity value of monitoring $VoI^A = VoI^M(1 + 1/r)$. The annuity value calculates the present value of a stream of income discounted at the rate r. Re-scaling the marginal value for one period VoI^M to the annuity value VoI^A allows a more direct comparison to the total value of monitoring and the net present value of the fishery, which are determined by the discounted stream of income from the fishery.

4.3.5.3. *Frequency of monitoring.* I calculated the frequency that the manager chose to monitor the stock under the optimal policy to characterize the importance of monitoring to the management strategy over the long run. These values are intended to guide long-run prioritization decisions like the target stock assessment frequency described in (**5**). I used simulations (appendix C.3) to compute how often the manager monitors the stock under the optimal policy.

4.3.5.4. *Global sensitivity analysis.* I used a global sensitivity analysis (GSA) to identify the effect of biological parameters and the management objective on the value of monitoring and frequency under the optimal policy. GSA quantifies the effect of varying each model parameter over a range of possible values on a chosen outcome. The parameters included in the sensitivity analysis and the respective ranges are given in table 4.1. I quantified the magnitude of the effect of each parameter using variance-based GSA (**14**). In addition, I characterized the sign and shape of the relationship with partial dependence plots (**21**).

GSA metrics require many repeated evaluations of the model and, therefore, are computationally expensive to compute directly. To mitigate this cost, I computed the GSA metric with an emulator, a machine-learning model approximating the true model that is cheap to evaluate repeatedly. I trained the emulators on a data set generated by calculating VoI^T, VoI^m, and $\omega_{i=1}$ at a set of input parameters sampled by a low discrepancy sequence (Sobol sequence) from the ranges used in the global sensitivity analysis. I then fit boosted regression tree models to each data set (**22**), and tested the quality of the approximations on a set of test data not included in the training process.

The correlation between the test data and the predicted values was 0.996 for the marginal value of monitoring, 0.996 for the total value, and 0.961 for the frequency of monitoring.

4.4. Results

4.4.1. Optimal monitoring strategy. The optimal choice to monitor the stock in a given period i_t depends on the current abundance estimate $\hat{B}_t$ and the associated level of uncertainty CV_t (the manager's belief state). The relationship between the belief state and the choice to monitor is called the policy function. The shape of the policy function in turn depends on the harvest control rule, management objectives, and biological characteristics of the stock. However, the policy functions do have some common features. In general, monitoring is optimal for large values of uncertainty, but the threshold value $\overline{CV}$ depends on the estimated biomass $\hat{B}_t$ (fig.4.4.A). The relationship between the estimated biomass and the threshold uncertainty is a decreasing function of the expected biomass, indicating that monitoring is more valuable when the stock is abundant (fig.4.4.B). When the harvest control rule includes a limit reference point, it is never optimal to monitor when the estimated abundance is small (fig. 4.4.A). Adding uncertainty buffers to the harvest control rule tends to increase the threshold uncertainty, indicating that monitoring is less valuable in these cases.

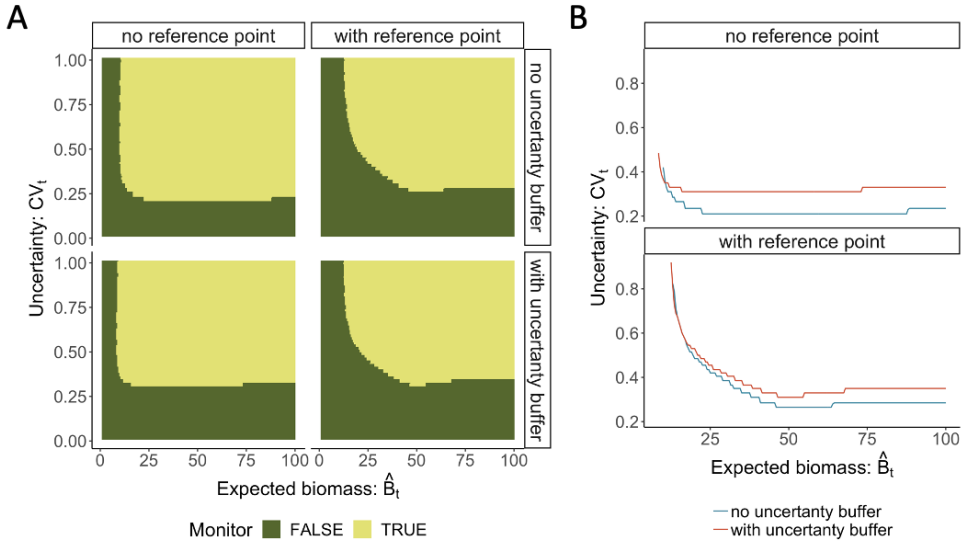

FIGURE 4.4. The policy functions A) heatmaps indicating regions of state space where monitoring is optimal (yellow) and where it is not (green). B) The threshold value of CV_t above which monitoring is optimal with and without buffers β added to account for uncertainty.

4.4.2. Economic value of monitoring. Monitoring programs can contribute a large fraction (¿20% in our base scenario) of the total economic value of a fishery. I characterize the economic value of monitoring information in two ways. The marginal benefit of a single observation is calculated as an annuity and the total net present value of the monitoring program. Both the marginal benefit and the total value of the monitoring programs depend on the estimated biomass $\hat{B}_t$ and level of uncertainty CV_t (fig. 4.5). The marginal benefit of monitoring is very sensitive to the current level of uncertainty and biomass varying between 0% and 60% in the base scenario. The total value of the monitoring program is less sensitive to the current belief state varying between 16% and 24% of the net present value of the fishery. Both values are larger in proportional terms when the harvest objective is considered in isolation, between 0% and 80% and 20% and 30%, respectively (fig. C.5). The increase in percentage terms can be attributed to the fact I demonstrate in the following section that non-consumptive values have a small effect on the value of monitoring but do increase the total value.

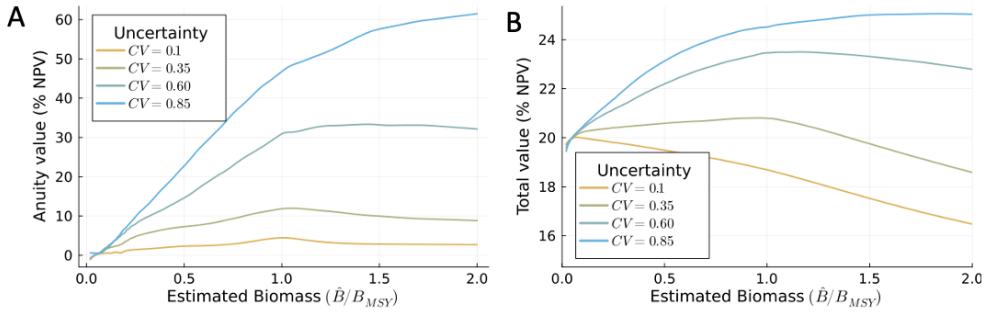

FIGURE 4.5. The economic value of the monitoring program in the base scenario. A) The marginal benefit of monitoring is calculated as an annuity, and B) the total value of the monitoring program as percentages of the net present value of the fishery. Not the difference in scale between the y axes

4.4.3. Effect of biological and management context. The economic value of monitoring programs was primarily derived from the fishery's profits, even though non-consumptive conservation goals were given equal weight in the manager's problem. The value of information was then modified by the biological characteristics of the stock and additional management context.

The marginal benefit of monitoring depends primarily on the current level of uncertainty CV_t and estimated biomass $\hat{B}_t$ (fig. 4.6.A). If the current level of uncertainty is low, then additional observation will have a limited effect. Similarly, when the stock is scarce, the choice in harvest

limit $\overline{H}_t$ is insensitive to small changes in abundance, limiting the value of new observations. In addition to the belief state and fishery profits, the size of the uncertainty buffer β, the quality of information in the absence of monitoring σ_p, and the growth rate of the stock F_{MSY} have a large positive effect on the marginal benefit of monitoring.

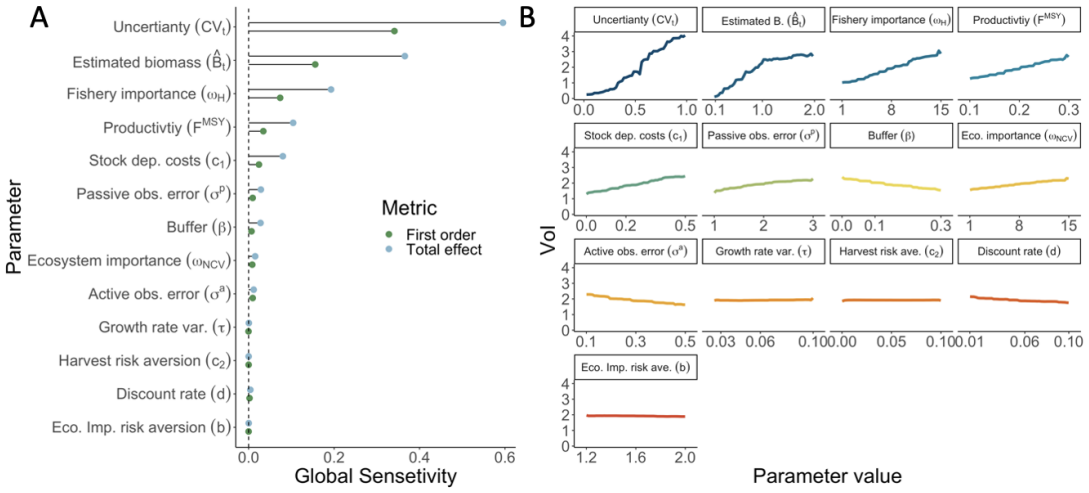

FIGURE 4.6. A) the global sensitivity of the marginal benefit of minoring VoI^M to each model parameter and the belief state variables. The global sensitivities are measured in two ways: the first-order effect, which only considers the effect of a change in the focal parameter, and the total effect, which includes interaction effects between multiple parameters. B) The partial dependence of the value of monitoring on each belief state variable and model parameters.

These same factors also contribute to the total value of the monitoring program (fig. C.6). However, we found that the belief state has a much smaller effect on the total value than the marginal benefit, while the discount factor had a larger influence on the total value. The total value of monitoring depends on the benefits from the program accumulating over a long time horizon, reducing the importance of the system's current state. The discount factor also becomes important because it determines the weight given to the value of observations made in the distant future.

The frequency of monitoring under the optimal policy was determined by many of the same factors that determined the economic value but was more sensitive to the growth rate of the stock F_{MSY} and the amount of environmental variability τ. The effect of environmental variability on monitoring is somewhat ambiguous ex-ante. A more variable environment increases the need for information because uncertainty grows more quickly over time, but one the other hand, any information you get from monitoring decays more quickly, decreasing its value. Here, we find that the first effect

seems to dominate, and environmental variability increases monitoring frequency. The large effect of the growth rate of the stock on monitoring frequency is caused by the fact that stocks with higher growth rates are harvested at higher target exploitation ratios F_t. Because the exploitation rate is higher, mistakes in setting the catch limit will compound more quickly over time.

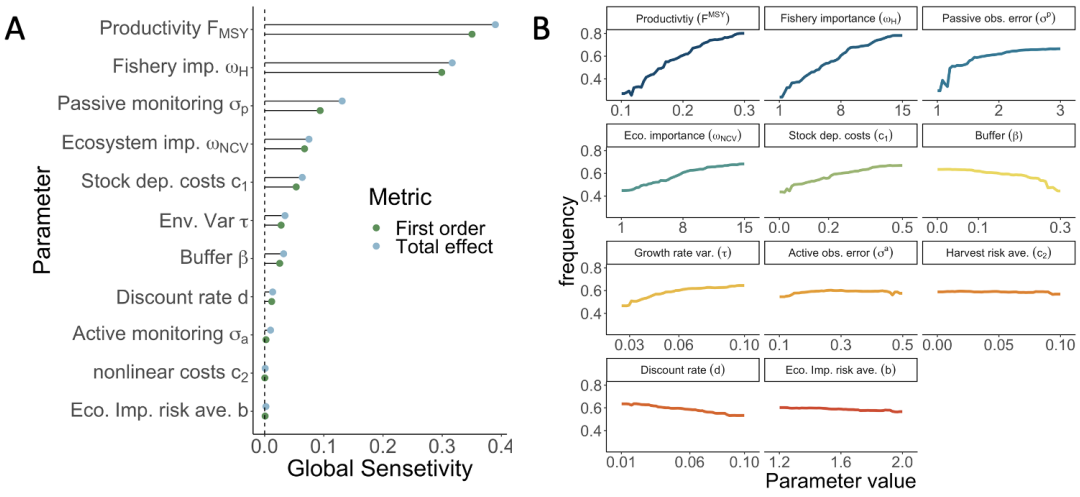

FIGURE 4.7. A) Importance of each model parameter for determining the frequency of monitoring under the optimal policy. B) The partial dependence of the frequency of monitoring on each model parameter.

4.4.4. Robustness tests. I tested the robustness of my findings to two key assumptions: 1) that the fishery cannot reduce the biomass of the stock below a minimum value B_{min} and 2) the uncertainty buffers β were constant over time. The qualitative results are generally robust to these assumptions (appendix C.1) with a few exceptions discussed below.

4.4.4.1. *Minimum biomass.* I found that the value of monitoring was largely unaffected by the choice in minimum biomass B_{min} except for high values of uncertainty (fig. C.1). When uncertainty is high there is a chance that the harvest limit $\overline{H}$ is set at or near the biomass of the stock B_t in which case the biomass limit B_{min} would bind. Small values of B_{min} increase the value of monitoring because these over-harvest events are more costly as it takes a longer amount of time for the stock to recover. Understanding the capacity of the fishery to deplete the stock is likely an important factor for designing programs and prioritizing monitoring resources for data-limited stocks.

4.4.4.2. *Uncertainty buffers ($p^* - \sigma$).* I tested the effect of buffers that scaled with the current level of uncertainty CV_t using the $p^* - \sigma$ rule (**20**). Ex-ante, The effect of these rules on the value of monitoring is ambiguous; it may increase the value of monitoring because greater uncertainty systematically reduces harvest limits or reduce its value because the probability of overfishing the stock because of uncertainty is reduced. I found the latter to be true. The marginal benefit and the total value of monitoring are lower when the $p^*\sigma$ rule is applied (fig.C.2). Furthermore, The frequency of monitoring under the optimal policy was reduced for larger uncertainty buffers (smaller values of p^*, fig.C.3). The $p^* - \sigma$ rule also changed the sensitivity of the optimal policy to the management objectives, increasing the effect of fishery profits ω_H and reducing the effect of conservation values ω_N (fig. C.3). The $p^* - \sigma$ rule reduces the effect of conservation values by reducing the risk of overharvest when uncertainty is high.

4.4.5. Comparison to U.S. National Marine Fisheries Service (NMFS) stock assessment prioritization framework. To illustrate how my findings can inform practical management decisions, I compared them to a framework developed by NMFS to prioritize model-based stock assessments. Stock assessments are used in commercial fisheries management to evaluate the abundance of harvest populations and identify sustainable harvest levels; thus they play an analogous role to the monitoring observations described by our model. Furthermore, stock assessments are updated at semi-regular intervals depending on the needs of the management agency. The NMFS framework discusses these prioritization decisions.

I found two primary differences between the NMFS prioritization framework and our model results (table 4.2): 1) the NMFS framework gives equal weight to ecosystem importance (non-consumptive values) and fishery importance (harvest value). In contrast, my results showed that harvest values are the primary driver of the value of the monitoring program. 2) The NMFS framework increases the priority of assessments for stocks that are below their target abundance. My model indicates that it is not valuable to monitor the stock at low abundances because the choice in harvest limit is not sensitive to new information under these circumstances.

In addition, the NMFS framework and my model include different biological characteristics of the stock. The NMFS framework uses the population age structure (mean and in catch) while my framework includes the stock productivity F_{MSY}. However, these two factors are often correlated; short-lived fast-growing species (younger mean age in catch) are consistently more productive than

Factor	NMFS framework	Model parameter	Annual priorities		Target frequency	
			NMFS	Model	NFMS	Model
Fishery importance	recreational, subsistence commercial	Harvest weight ω_H	$0-15$	0.075	+1 yr	0.3
Non-consumptive values	ecosystem importance	NCV weight ω_{NCV}	$0-10$	0.008	+1 yr	0.06
Stock status	Abundance, rebuilding status	Estimated biomass $\hat{B}_t$	$0-6$	0.15	0	0
	relative fishing mortality	Not considered	$0-5$	0	0	0
Assessment information	years assessment overdue	Uncertainty CV_t	$0-10$	0.34	0	0
	Assessment level	Measurement accuracy σ^a	0	0.009	0	0.002
	New data, changes in indicators	Not considered	$0-10$	0	0	0
	Not considered	passive info. σ^p	0	0.009	0	0.09
Stock biology	mean age in catch	Not considered	0	0	$\times 1/2$	0
	Not considered	Productivity F_{MSY}	0	0.035	0	0.35
	Recruitment variability	productivity shocksτ	0	0.00	$\pm$1 yr	0.03
Policy process	Constituent demand	Not considered	$0-5$	0	0	0
Harvest Control rule	Uncertainty buffer	Uncertainty buffer β	0	0.006	0	0.025

TABLE 4.2. Factors included in the NMFS stock assessment prioritization framework (**5**) and the corresponding factor included in the model. Factors that are not included in the framework have a weight of "0" and while factors that are tested and have no weights are marked "0.00"

long-live species (**23**). My finding that more productive stocks are assessed more frequently under the optimal policy is consistent with the NMFS framework, which increases assessment frequency for stocks with lower mean ages in catch.

4.5. Discussion

My analysis demonstrates that monitoring programs can significantly increase the economic value of a fishery, with contributions ranging from 5% to 20% in the total value in our base scenario. This value stems from improved harvest decisions, which allow for higher harvest limits for a risk of over fishing the stock. As such, the value of monitoring is primarily derived from fisheries profits rather than non-consumptive conservation values (ecosystem importance). This does not mean that conservation cannot justify monitoring in general because this information could be useful for other conservation measures, like identifying critical habitats not considered by our model.

Management context and biological factors also influenced the value of monitoring. The availability of alternative information sources reduced the value of additional investments in monitoring; this is described by the passive monitoring variable σ_p in our model. This parameter is important for scenarios where multiple information sources are available. For example, many fisheries management systems combine fishery-dependent data and independent surveys to estimate the stock biomass. In this context, the value of the independent surveys depends on the quality of the fishery-dependent data. The optimal monitoring strategy was also effected by adding buffers to reduce the harvest limits; this suggests a difference between the costs of reducing harvest with the buffers and the cost of additional monitoring. Finally, the growth rate of the stock F_{MSY} increased the value of monitoring programs. Indicating that faster-growing, short-lived species require more monitoring resources.

4.5.1. Mechanistic interpretations. In addition to providing a quantitative estimate of the economic value of monitoring programs, my findings illustrate two general principles: 1) information is valuable only if it is likely to change a decision (**2**, **24**) and 2) monitoring (and the associated learning) helps establish stabilizing feedback mechanisms (*sensu.* (**3**)). The first principle explains the sensitivity of the optimal monitoring policy to the choice in harvest control rule. Monitoring was most valuable when the expected stock size was larger than management targets. In this region of the state space, changes in estimated stock size result in relatively large changes in the harvest limit. When the stock is below management targets, the harvest limit is much less sensitive to new information and, therefore, less valuable. These findings align with previous research on optimal monitoring strategies, which demonstrated that monitoring is most beneficial near thresholds where

discrete changes in management actions occur (**12**, **13**). However, in this case, the management action is continuous, so the value of monitoring is increased over a large region of the state space where the harvest limits are sensitive to the estimated biomass.

4.5.2. Monitoring as a stabilizing feedback. The second principle links the economic value of monitoring to its effects on the dynamics of the coupled human–natural system. Accurate and frequently updated monitoring information allows managers to respond to opportunities created by increases in abundance and avoid overfishing the stock if it declines in abundance. This responsiveness is especially important for setting fisheries catch limits because it eliminates a positive feedback mechanism that can create a tipping point and cause fishery to collapse (**7**). Consider an extreme case where monitoring is absent, and the harvest levels are constant over time. The amount of biomass that can be removed sustainably from a population is determined by its abundance; therefore, if the stock declines in abundance, the sustainable harvest level might also decline. If the harvest limits remain constant, they can exceed the new lower sustainable level, causing further declines in abundance and initiating a vicious cycle. Preventing this vicious cycle may explain the relatively high value of monitoring in this system.

This second principle can also explain why monitoring is more valuable for stocks with higher productivity. We found that more productive stocks (higher F_{MSY}) were monitored more frequently under the optimal policy, indicating that tighter feedback between abundance and harvest limit is valuable in these scenarios. One explanation is more productive stocks are exploited at higher ratios (a larger fraction of the stock is harvested each period) and, therefore, can become overfished more quickly if the stock declines and corrective actions are not taken. This mechanism might also explain why monitoring is more valuable when the harvest control rule has a smaller uncertainty buffer (higher values of β), which also leads to higher exploitation ratios. A second explanation is the effect of productivity on the time scale of the population dynamics. Larger values of F_{MSY} are associated with fast-growing, short-lived species (**23**) that have faster population turnover and whose abundance varies on shorter time scales. In this context, more frequent monitoring is needed to match the faster time scale of the population dynamics.

The value of monitoring for establishing stabilizing feedback mechanisms also points to a limitation of our modeling framework: the absence of alternative feedback mechanisms. In the absence of strong regulations, bioeconomic theory suggests that fisheries are self-stabilizing. All else equal, as

stocks decline, fish become more difficult to find and expensive to harvest until an equilibrium is reached where harvest costs are equal to revenues (**25**, **26**), albeit at a lower equilibrium stock and profits relative to a managed fishery. This process creates a stabilizing feedback mechanism that can reduce the extent of overfishing if managers fail to adjust catch limits appropriately. My model does not incorporate these alternative stabilizing mechanisms, which could cause it to overestimate the value of the monitoring program. However, this stabilizing feedback can be weakened if the price of catch greatly outweighs the costs of harvesting, the stock remains easy to find when rare (**27**,**28**,**29**), or the stock is exploited as part of a multi-species fishery (**30**). Although including them would likely increase the value of monitoring, these additional economic factors are beyond the scope of this study.

4.5.3. Incentives to invest in monitoring. In addition to identifying factors that determine the value of information, the model also illustrates where the incentives to invest in monitoring come from. Specifically, we find that the value of monitoring is largely derived from fisheries profits. This finding aligns with previous results on optimal harvest strategies, demonstrating that under a wide range of conditions, a sole owner of a fishery has a strong incentive to conserve the stock (**26**, **31**). Under these conditions, the sole owner limits harvest in the present, given their understanding of how it affects future profitability. Similarly, the optimal monitoring decision takes into account how monitoring costs today influence future profits. However, if the profits from the fishery are low relative to the cost of monitoring, this incentive might not be strong enough to justify the investment, which can compromise sustainability.

4.5.4. Implications for designing harvest control rules. My findings also have implications for designing fisheries' harvest control rules. Multiple recent studies have identified modifications to HCRs as a simple and potentially effective reform to increase the resilience of fisheries to climate change (**8**, **32**, **33**). My results indicate that these reforms will also influence the optimal strategy for investing in monitoring information. Specifically, I find that adding uncertainty buffers and reference points to HCRs makes management less information-hungery; especially when stocks are scarce. This is logical because both measures are designed to improve the robustness of the harvest strategy to uncertainty (**33**). Reducing the information requirements of harvest strategy may be a valuable contribution of these reforms, especially under climate change, where increasing

variability in the abundance and distribution of stocks could make reliable abundance estimates more costly; on the flip side, where robust monitoring systems are in place changes to HCRs maybe less valuable.

4.5.5. Model assumptions and limitations. Although these results have clear implications for investing in monitoring information and designing harvest strategies, they need to be considered in light of simplifying assumptions and limitations of the model. Critically, we assume that the catch limits chosen by the manager will be a binding constraint unless the stock's abundance is very low $B_t < 5\%B_{MSY}$. This may be a reasonable assumption for highly valuable commercial or popular recreational fisheries, but it limits the applicability of our results to less valued species where the harvest limits do not constrain harvest levels. This assumption is useful because it allows the model to take as given the decisions made by harvesters, simplifying our analysis. In practice, however, the adaptive behavior of these decision-makers may need to be considered. For example, increases in fishing efforts might signal a need for additional monitoring.

The model also uses a simplified representation of uncertainty. First, I assume the manager knows the exact functional form and parameter value for the stock's population dynamics. Relaxing this assumption would cause uncertainty to increase more rapidly between monitoring observations, likely increasing the frequency of monitoring under the optimal policy. The model's representation of state uncertainty (measurement error) is also simplified. Abundance estimates from stock assessments typically have two sources of uncertainty: an estimate of abundance relative to a historical baseline and a scaling factor relating changes in relative abundance to changes in biomass (**34**). Both factors are important for implementing harvest control rules. Monitoring changes in relative abundance allows harvest levels to respond to changes in abundance, establishing a stabilizing feedback mechanism, while the scaling factor is used to calibrate the harvest control rule (**5**) (i.e. how much to increase or decrease harvest in response to a change in relative abundance). My model assumes the HCR is well-calibrated and focuses on uncertainty in relative abundance. Therefore, it does not capture the value of monitoring for improving estimates of the scaling factor and may overestimate the performance of the harvest control rules.

More broadly, my analysis is based on a decision-theoretic approach that weighs the costs and benefits of monitoring to inform a specific decision (catch limits), assuming the manager follows a

predetermined rule (the HCR). By focusing on one specific decision, the model may tend to underestimate the value of monitoring information. The model ignores the value of monitoring information for decisions like gear or area restrictions for the fishery. It also ignores the value of transferable knowledge produced by long-term monitoring data sets (**35**). In contrast, my assumption that the manager will strictly follow a predetermined rule will cause the model to overestimate the value of information because factors like status-quo bias and conflict between interest groups might limit the impact of new information on management actions, reducing its potential value (**35**).

4.6. Conclusion

Fisheries are an example of coupled human and natural systems where human activities and the state ecosystem are linked in a cycle of mutual causation. Fisheries management, defined by a combination of data collection schemes, decision-making rules, and regulatory enforcement (**32**), can be a stabilizing feedback mechanism linking ecosystem states and the level of human impacts. Tightening these feedback loops can increase the resilience of the coupled system, improving economic and environmental outcomes (**3**). Monitoring information can be costly but forms a critical link in the cycle. We find that efficient strategies for investing in monitoring depend on the choice harvest control rule, the stock biology, and the management objectives. This suite of factors can help regulators prioritize scarce scientific resources, improving the efficiency and resilience of fisheries and the ecosystem services they provide.

Bibliography

1. C. J. Walters, C. S. Holling, *Ecology* **71**, 2060 (1990).
2. E. McDonald-Madden, I. Chadès, M. A. McCarthy, M. Linkie, H. P. Possingham, *Ecological Applications* **21**, 844 (2011).
3. S. A. Levin, J. Lubchenco, *BioScience* **58**, 27 (2008).
4. C. J. Legg, L. Nagy, *Journal of Environmental Management* **78**, 194 (2006).
5. R. Methot, *U.S. Dep. Commer., NOAA Tech. Memo. NMFS-F/SPO152, 31 p* p. 44 (2015).
6. C. Walters, P. H. Pearse, *Reviews in Fish Biology and Fisheries* **6**, 21 (1996).
7. J. Roughgarden, F. Smith, *Proceedings of the National Academy of Sciences* **93**, 5078 (1996).
8. J. P. Kritzer, C. Costello, T. Mangin, S. L. Smith, *ICES Journal of Marine Science* **76**, 1424 (2019).
9. S. J. Newman, *et al.*, *Marine Policy* **88**, 11 (2018).
10. *The State of World Fisheries and Aquaculture 2020* (FAO, 2020).
11. M. Memarzadeh, C. Boettiger, *The American Naturalist* **193**, 645 (2019).
12. D. M. Kling, J. N. Sanchirico, P. L. Fackler, *Journal of Environmental Economics and Management* **84**, 223 (2017).
13. M. R. Sloggy, D. M. Kling, A. J. Plantinga, *Journal of Environmental Economics and Management* **103**, 102357 (2020).
14. I. Sobol, *Mathematics and Computers in Simulation* **55**, 271 (2001).
15. A. E. Punt, D. S. Butterworth, C. L. de Moor, J. A. A. De Oliveira, M. Haddon, *Fish and Fisheries* **17**, 303 (2016).
16. O. Le Pape, S. Bonhommeau, A.-E. Nieblas, J.-M. Fromentin, *Proceedings of the National Academy of Sciences* **114** (2017).
17. M. Springborn, J. N. Sanchirico, *Journal of Environmental Economics and Management* **66**, 609 (2013).
18. C. M. Free, *et al.*, *Fish and Fisheries* **24**, 248 (2023).
19. P. L. Fackler, R. G. Haight, *Resource and Energy Economics* **37**, 226 (2014).
20. M. H. Prager, C. E. Porch, K. W. Shertzer, J. F. Caddy, *North American Journal of Fisheries Management* **23**, 349 (2003).
21. B. M. Greenwell, *The R Journal* **9**, 421 (2017).
22. B. Greenwell, B. Boehmke, J. Cunningham, G. Developers, *gbm: Generalized Boosted Regression Models* (2020). R package version 2.1.8.

23. S. Zhou, S. Yin, J. Thorson, A. Smith, M. Fuller, *Canadian Journal of Fisheries and Aquatic Sciences* **69**, 1292 (2012).
24. S. Canessa, *et al.*, *Methods in Ecology and Evolution* **6**, 1219 (2015).
25. V. L. Smith, *Journal of Political Economy* **77**, 181 (1969).
26. C. Clark, *Mathematical Bioeconomics: The Optimal Management of Renewable Resources.* (Wiley, 1993).
27. A. S. Golden, B. van Poorten, O. P. Jensen, *Fish and Fisheries* **23**, 1418 (2022).
28. B. E. Erisman, *et al.*, *Canadian Journal of Fisheries and Aquatic Sciences* **68**, 1705 (2011).
29. C. J. Dassow, *et al.*, *Canadian Journal of Fisheries and Aquatic Sciences* **77**, 762 (2020).
30. M. G. Burgess, S. Polasky, D. Tilman, *Proceedings of the National Academy of Sciences* **110**, 15943 (2013).
31. W. J. Reed, *Journal of Environmental Economics and Management* **6**, 350 (1979).
32. A. E. Punt, *et al.*, *ICES Journal of Marine Science* **71**, 2208 (2014).
33. T. K. Mildenberger, *et al.*, *Fish and Fisheries* **23**, 73 (2022).
34. R. Hilborn, C. J. Walters, *Quantitative fisheries stock assessment: choice, dynamics and uncertainty* (Springer Science & Business Media, 2013).
35. J. N. Sanchirico, M. R. Springborn, M. W. Schwartz, A. N. Doerr, *Conservation Biology* **28**, 361 (2014).

APPENDIX A

Vaccine Prioritization

A.1. Model specification, parameterization and optimization

A.1.1. Model dynamic equations. The dynamic equations specifying transitions between the disease states are as follows:

$$\dot{S}_i = -qs_i\theta \left[\sum_{j\in J} \sum_{m\in M} \tau_m r_{m,i,j} S_i \frac{I_{m,j}}{N_j} \right] - \mu_i v \tag{A.1}$$

$$\dot{F}_i = -qs_i\theta \left[\sum_{j\in J} \sum_{m\in M} \tau_m r_{m,i,j} F_i \frac{I_{m,j}}{N_j} \right] + (1-\epsilon_i)\mu_i v \tag{A.2}$$

$$\dot{E}_i = qs_i\theta \left[\sum_{j\in J} \sum_{m\in M} \tau_m r_{m,i,j} (S_i + F_i) \frac{I_{m,j}}{N_j} \right] - E_i/\gamma_{exp} \tag{A.3}$$

$$\dot{P}_i = \epsilon_i \mu_i v \tag{A.4}$$

$$\dot{I}_{pre,i} = E_i/\gamma_{exp} - I_{pre,i}/\gamma_{pre} \tag{A.5}$$

$$\dot{I}_{asym,i} = \sigma_{asym} I_{pre,i}/\gamma_{pre} - I_{asym,i}/\gamma_{asym} \tag{A.6}$$

$$\dot{I}_{sym,i} = (1-\sigma_{asym}) I_{pre,i}/\gamma_{pre} - I_{sym,i}/\gamma_{sym} \tag{A.7}$$

$$\dot{R}_i = I_{asym,i}/\gamma_{asym} + (1-\delta_i) I_{sym,i}/\gamma_{sym} \tag{A.8}$$

$$\dot{\gamma}_i = \delta_i I_{sym,i}/\gamma_{sym}. \tag{A.9}$$

To reduce clutter we have suppressed the time index t on each of the state variables, the vaccine allocation vector μ_i, and the vaccination rate v.

Parameter	Description	Base Value(s)	Source
J	demographic groups: (1) age-only, (2) age and essential workers	{0-4, 5-19, 20-39, 40-59, 60-74, 75+}, {0-4, 5-19, 20-39, 20-39*, 40-59, 40-59*, 60-74, 75+}	Assumed
σ_{asym}	infection asymptomatic rate	0.16	(**1**)
δ	infection fatality rate (age-specific)	{6.7×10^{-6}, 2.5×10^{-5}, 0.0002, 0.002, 0.018, 0.12}	(**2**)
s	susceptibility (age-specific)	{0.5, 0.5, 1.0, 1.0, 1.0, 1.0}	(**3**)
τ_{pre} τ_{asym} τ_{sym}	relative infectiousness by symptom type	0.51 0.51 1.0	(**4**)
γ_{exp} γ_{pre} γ_{asym} γ_{sym}	symptom duration (days)	3.0 3.2 3.5 7.0	(**4**)
ϵ	vaccine effectiveness	0.9	Informed by initial COVID-19 vaccine effectiveness estimates (**5**)
p	proportion of essential workers	0.40	Calculated with labor data (**6**,**7**); alternative: (**8**)
R_0	secondary infections in a naive population	2.5	(**9**), (**10**)
q	transmission probability in a naive population	0.053	Calculated given R_0, s and other parameters
θ	scaling factor for transmission probability due to NPI other than social distancing	0.65	Assumed (consistent with estimated COVID-19 R_0 under NPIs (**10**))
n	population shares: (1) age-only, (2) age and essential workers	{0.06, 0.19, 0.27, 0.26, 0.19, 0.04}, {0.06, 0.19, 0.19, 0.08, 0.18, 0.8, 0.19, 0.04}	(**11**)
e	remaining years of life expectancy (age-specific)	{76, 66, 50, 31, 17, 6}	(**12**)
$R(0)$ $I_{sym}(0) + I_{asym}(0)$	initial recovered initial sympt. and asympt.	0.08 0.005	Informed by IHME projections (**13**)
v	fraction of population vaccinated daily	0.1/30	Informed by comments from CDC Director to U.S. Senate Panel (**14**)

TABLE A.1. Base model parameter values and sources.

A.1.2. Model parameters.

Scenario	Change from Base scenario parameters	Source
High initial infections	15 symptomatic infections per 1000	Assumed
Strong NPI	$\theta = 0.5$	Assumed
Weak NPI	$\theta = 0.75$	consistent with 30-70% of U.S. population always wearing a mask (**15**) with 33-58% effectiveness (**16**)
Weak vaccine	$\epsilon_i \in \{0.5, 0.5, 0.5, 0.5, 0.5\}$	Minimum value from FDA approval
Weak vaccine seniors	$\epsilon_i \in \{0.9, 0.9, 0.9, 0.9, 0.5, 0.5\}$	Informed by influenza vaccine effectiveness
High susceptibility ages < 20	$s_i \in \{1.0, 1.0, 1.0, 1.0, 1.0\}$	Assumed
Low supply	$v(t) = 0.05/30$ allocation policies switched every 60 days	Assumed
Ramp up	$v(t) = \left\{ \begin{array}{ll} 0.05/30, & t \leq 60 \\ 0.10/30, & t > 60 \end{array} \right]$ Allocation polices switch for every 10% of population vaccinated	Informed by comments from the scientific head of the U.S. vaccine development program (**17**)
Open schools	$\alpha_{school} = 0.7$	Assumed
High contacts	$\alpha_{social} = 0.5$	Assumed

TABLE A.2. Parameter values that differ from the Base case for alternative scenarios.

A.1.3. Parameters for alternative scenarios.

A.1.4. Initial conditions. The number of susceptible, infected and recovered individuals is likely to vary by region and will depend on the time when the vaccine is deployed. Because of the likely variation in this parameter, we test a range of values from 1-20 symptomatic cases per 1000 (when the vaccine is deployed). The infections are assumed to be distributed between groups in accordance with the stable distribution of cases when the epidemic is growing exponentially. The portion of each group infected at time $t = 0$ in the Base parameter set is shown in Fig. A.1. The number of recovered individuals in the population was set to 8% of each demographic group which was informed by IHME projections (**13**). When the vaccine is first deployed, we set the number of deceased and vaccinated individuals in each age group to zero (since the share of actual deceased is very small, i.e. approximately 0.0015). The number susceptible was set to ensure the proportion of individuals in each category summed to 100% after the number of individuals at each stage of infection was determined.

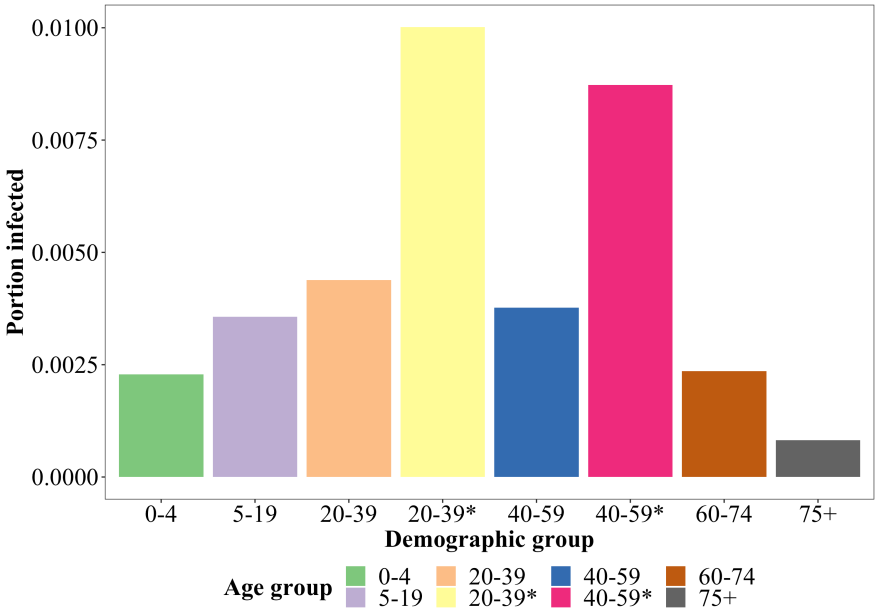

FIGURE A.1. The proportion of each demographic group infected when vaccine distribution begins in the Base scenario.

A.1.5. Calibration. The relationship between the basic reproduction number, R_0, and parameters governing transmission and epidemiological characteristics is given by the so-called next-generation matrix:

$$R_0 = max[eigenval\{q\gamma_{pre}\tau_{pre}(r\Delta s)\Delta n + q\sigma_{asym}\gamma_{asym}\tau_{asym}(r\Delta s)\Delta n + q(1-\sigma_{asym})\gamma_{sym}\tau_{sym}(r\Delta s)\Delta n\}], \tag{A.10}$$

where the maximum eigenvalue operator wraps several terms including r, the social contact matrix, s, the age-specific susceptibility rate, n, a vector of the proportions of the population in each demographic group and Δ, an operator that signifies multiplying each row of a matrix by the corresponding entry in the vector. For symptom type $m \in \{asym, pre, sym\}$, the constants γ_m, τ_m and σ_m represent the duration, relative infectiousness of an individual and the probability of type m, respectively.

We first set a baseline $R_0 = 2.5$ as estimated by (**9**). We then solve for the transmission probability parameter, q, using Equation A.10, assuming a naive (pre-pandemic) population. We then scaled q by a fixed factor $\theta \in [0, 1]$ to reflect the impact of non-pharmaceutical interventions (NPI) like masks, hand washing and maintaining distance when contacts are made.

A.1.6. Contact matrices distinguishing essential workers. Estimated contact rates for the U.S. were obtained from (**18**) who used population-based contact diaries from the European POLYMOD survey to project to other countries, including the U.S. These included contact rates for 16 age classes in five year increments from ages 0 to 80. We collapsed these to five age groups (0-4, 5-19, 20-39, 40-59, 60-80) using population-weighted sums:

$$\hat{r}_{i,j,x} = \sum_{\mathtt{i}\in i} \left(\frac{N^{pop}_{\mathtt{i}}}{N^{pop}_{i}} \sum_{\mathtt{j}\in j} r_{\mathtt{i},\mathtt{j},x} \right), \tag{A.11}$$

where $\{\mathtt{i},\mathtt{j}\}$ are the subscripts for the five year age bins, $\{i,j\}$ are the subscripts for the larger age bins, $r_{\mathtt{i},\mathtt{j},x}$ is the average number of daily contacts a person in group $\mathtt{i}$ makes with a person in group $\mathtt{j}$ for activity/location x, and N^{pop}_{h} is the population size for age group h.

The total number of i-to-j contacts must equal the total number of j-to-i contacts: $N^{pop}_{i} r_{i,j,x} = N^{pop}_{j} r_{j,i,x}$. Because numerical issues—estimation in (**18**), bin discretization and rounding—can lead to small differences, we ensure this condition holds by imposing,

$$r_{i,j,x} = \frac{0.5(N^{pop}_{i}\hat{r}_{i,j,x} + N^{pop}_{j}\hat{r}_{j,i,x})}{N^{pop}_{i}}, \tag{A.12}$$

where the numerator is the mean of the two measures of total contacts between groups i and j and the denominator transforms the result to per-capita in i.

Setting essential worker contact rates requires additional assumptions and attention to the activity/location. We define the essential worker indicators $e \in \{n, y\}$ for "no" and "yes". Our grouping is such that all essential workers ($e = y$) are employed but non-essential-workers ($e = n$) are a mix of employed and not employed. Let e' represent the indicator for a second group which can be equal or not equal to the value for e.

In the case of all activities/locations x that are not $work$, contact rates are given by

$$r_{(i,e),(j,e'),x} = \frac{N^{pop}_{(j,e')}}{N^{pop}_{j}} r_{i,j,x}, \;\; \forall\, x \neq work. \tag{A.13}$$

This follows from the assumption that contacts made by any group (i,e) with any other group (j,e') are independent of i's essential worker status. Thus, we only need to split contacts $r_{i,j,x}$ into those made with essential worker type $e' = y$ versus the remainder with type $e' = n$, i.e. given the share $N^{pop}_{(j,e')}/N^{pop}_{j}$.

Estimating contacts when $x = work$ involves a larger number of steps. We first address contacts made by essential workers ($e = y$) before turning to non-essential workers ($e = n$). For $e = y$, let the share of the working age population $(20 - 59)$ in group i that is employed be given by p_i.

We assume that all of the work contacts are attributable to employed adults resulting in an employed adult contact rate of $r_{i,j,work}/p_i$. Then the contact rate of essential workers ($e = y$) in group i with age group j is

$$r_{(i,y),j,work} = \frac{r_{i,j,work}}{p_i}. \tag{A.14}$$

Let the fraction of working age group i that is employed in an essential worker role be given by $p_{i,y}$. The average workplace contact rate for non-essential-workers in group i with group j is given by

$$r_{(i,n),j,work} = \alpha_{work} \left(\frac{r_{i,j,work}}{p_i} \right) \left(\frac{p_i - p_{i,y}}{1 - p_{i,y}} \right), \tag{A.15}$$

where $\alpha_{work} < 1$ scales for social distancing and the final term in brackets scales for the share of non-essential workers that are employed and thus have contacts at $work$.

Finally, we assume that the average $work$place contact rate for an individual of type (i, e) with individuals of type (j, e') is given by the partial contact rate $r_{(i,e),j,work}$ times the proportion of total work contacts of individuals in group j that are made by individuals in sub group e':

$$r_{(i,e),(j,e'),work} = r_{(i,e),j,work} \left(\frac{N_{(j,e')} \cdot r_{i,(j,e'),work}}{N_{(j,y)} \cdot r_{i,(j,y),work} + N_{(j,n)} \cdot r_{i,(j,n),work}} \right). \tag{A.16}$$

In addition to The formulation described above we also considered a scenario where essential workers contacts were clustered (i.e. individuals only contact others of the same essential worker status at work). The work contact rates for each group are calculated in the same manner as described above, but we assume that the c work contacts between essential and non-essential workers are zero.

These two model formulation represent two extremes. Work contacts are likely to be concentrated among others of the same essential status (as opposed to formulation one) but essential workers are likely to have some contacts non-essential workers in the work place.

Finally we scale the work contacts for age groups that are not separated into essential and non-essential workers (5-19, 60-80) to match the scaling for prime working age classes.

$$r_{i,j,work} = \left(\frac{p_{i,y}}{p_{work}} + \alpha \frac{p_{work} - p_{i,y}}{p_{work}} \right) r_{i,j,work} \tag{A.17}$$

A.1.7. Optimization algorithm. The optimization algorithm used in our analysis is split into two parts. First a genetic algorithm is run to identify an effective strategy near a global optimum. This solution is then refined via simulated annealing.

Genetic algorithms take inspiration from the natural process of evolution, and work by randomly sampling a populating of candidate solutions, selecting a set of survivors based on the candidates performance against the objective function, information from these survivors is then used to generate a new generation of candidates solutions, and so forth (**19**). The genetic algorithm executes the following steps:

(1) Sample $N_{t=0}$ candidate solutions $\{x_{n,t=0}\}$ from a Dirichlet distribution with parameter α_0.
(2) Each candidate solution is evaluated with the objective function.
(3) The bests $K_{t=0}$ candidates $\{x^{best}_{n,t=0}\}$ are solved and the distributions parameter α_0 is updated to α_1 which is the mean of $\{x^{best}_{n,t=0}\}$ times the entropy parameter for that time step η_t. The entropy parameter determines how concentrated new samples will be around the mean of the selected samples in the prior step.
(4) Steps 1 to 3 are repeated for a fixed number of iterations T and the best candidate solution sampled at any iteration is returned. The values N_t, K_t and η_t are tuned for each step to maximize performance.

Simulated annealing is based on thermodynamic models of cooling metals. Briefly, the algorithm is initialized by sampling a candidate solution x_0, this candidate solution is updated by sampling a new candidate solution x_t from a proposal distribution centered around x_0. This solution is either accepted and replaces the current x_0 or it is rejected and a new candidate solution is drawn using the existing value of x_0. The proposed solutions x_t are accepted if they perform better against the objective than the incumbent x_0, if $x_t > x_0$ it is selected with probability $p = exp\left[-(x_t - x_0)/T_t\right]$. large values of T_t increase the probability that a new candidate solution will be accepted allowing the algorithm to explore the solution space and move away from local minima. T_t is reduced over

time to allow the algorithm to start exploring the solution space and then eventually stabilize on a global minimum. The simulated annealing executes the following steps:

(1) Initialize a chain with value x_0. Generate a new sample from the proposal distribution $x_t \sim tr(N(tr^{-1}(x_0), \sigma I))$ where the transform tr from $\mathbf{R}^n$ to the solution space. Initialize a counter i that tracks the number of iterations.

(2) If $x_t < x_0$ replace x_0 with x_t, update $i = i + 1$ and repeat from step 1.

(3) If $x_t > x_0$ sample $\mu \sim unif(0, 1)$. If $\mu > exp\left(-(x_t - x_0)/T(i)\right)$ then replace x_0 with x_t update i and repeat from step 1. Otherwise save x_0 and repeat from step 1. We used $T(i) = T_0/i$ as the temperature function.

(4) Stop when $i > max\_iter$

These algorithms were tuned experimentally to consistently converge on a minimum solution on a test case. We used the minimum years of life lost under the Base parameter set as our test case. We found that numerical errors (defined as difference between the test runs) increased over the decision periods. This is caused by the face that the final decision periods are a region of very flat payoff, leading to a large number of solutions that perform very similarly but differ in this region To quantify the sensitivity of the solutions to deviations in the outcome of interest, for each decision variable, we identified the range of alternative values nearby that produced a similar objective function value. This procedure is described in Appendix A.1.8.

A.1.8. Whiskers on optimal vaccine allocation bars in Fig. 2.2. The whiskers on optimal vaccine allocation bars in Fig. 2.2 show the range of alternative allocations that still produce an outcome that is within 0.5% of the optimum. The upper (lower) bound of each whisker was produced one at a time by systematically exploring higher (lower) levels of the given decision variable (proportion of vaccines allocated to a given demographic group in a given decision period). Let x represent the level of a single decision variable. The whiskers for x were found in a two step process. First, approximate upper (lower) bounds on the range of x were found by sampling new candidate solutions above (below) the optimum with a Markov chain using the algorithm below.

(1) Initialize a chain with the optimized value $x_0 = x^*$. Generate a new sample from the proposal distribution $tr(x_t) \sim N(tr^{-1}(x_0), \sigma I)$, where the function applies the soft max function to each decision period so that the solutions are represented by a vector in

$R^{(n_{groups}-1)*n_{steps}}$, I is the identity matrix and $\sigma = 0.001$. $n_{groups} = 8$ is the number of demographic groups and $n_{steps} = 6$ is the number of decision periods. Initialize a counter i that tracks the number of iterations.

(2) If $f(x_i) < 1.005 * f(x^*)$, where f returns the value of the objective function, replace x_{i-1} with x_i, update $i = i + 1$, save x_i and sample a new candidate solution centered at x_i.

(3) Else reject x_i and sample a new candidate solution centered at x_{i-1}.

(4) Stop when $i > N_{samples} = 10000$

(5) Repeat steps 1-4 to generate samples from $N_{chains} = 30$ independent chains. For each decision variable x, return the maximum ($\overline{x}$) and minimum ($\underline{x}$) value from the list selected samples.

For each decision variable, this algorithm produces the pair $(\underline{x}, \overline{x})$, an upper bound for the minimum whisker value and, similarly, a lower bound for the maximum value for the whisker. We refined this solution as follows. For each decision variable, we take $\underline{x}$ and explore alternative candidate values $x^{cand} \in [0, \underline{x}]$, i.e., values between the lower bound on the decision variable and $\underline{x}$. To do so we use the bisection root finding algorithm adjusting the value of x^{cand} until $f(x^{cand}) - 1.005 * f(x^*) \leq tol = 0.01$. The procedure is repeated for the upper extent of the whisker for $x^{cand} \in [\overline{x}, 1]$. Thus the final whisker extents are each within a percentage point of the true bound.

A.2. Static policies

We found that the patterns observed for our dynamic solutions were similar for the static policies. The static policies targeted the same high priority groups: ages 5-19 and essential workers when minimizing infections, and ages 60+ along with essential workers when minimizing deaths and YLL. Again consistent with the dynamic policies we found that optimal vaccine allocations were substituted towards essential workers when the effectiveness of the vaccine was low and when the reproductive number when the vaccine first became available was small. These patterns are shown in Fig. A.2A. [1]

In Fig. A.2B we show how robust the static policies are when applied to the "wrong" scenario. As with the dynamic policies we found that most polices performed very poorly when applied to the weak vaccine ages 60+ and the strong NPI scenarios, relative to the optimum. But, in general,

[1]It should be noted that the quantities plotted in Fig.A.2A correspond to the allocation of the initial supply, which is different from the similar main text Fig.2.4 that presents the percent of each group vaccinated at 3 months.

deterioration in performance (due to a mismatch between the true scenario and the one driving the policy applied) was much worse for static than dynamic polices. One driver of this effect is that static polices identify one set of high priority groups and do not switch. Dynamic polices also identify one set of high priority groups but then also switch. Thus dynamic policies differ by when a group is prioritized, as opposed to the static polices which differ by which groups are a high (unchanging) priority.

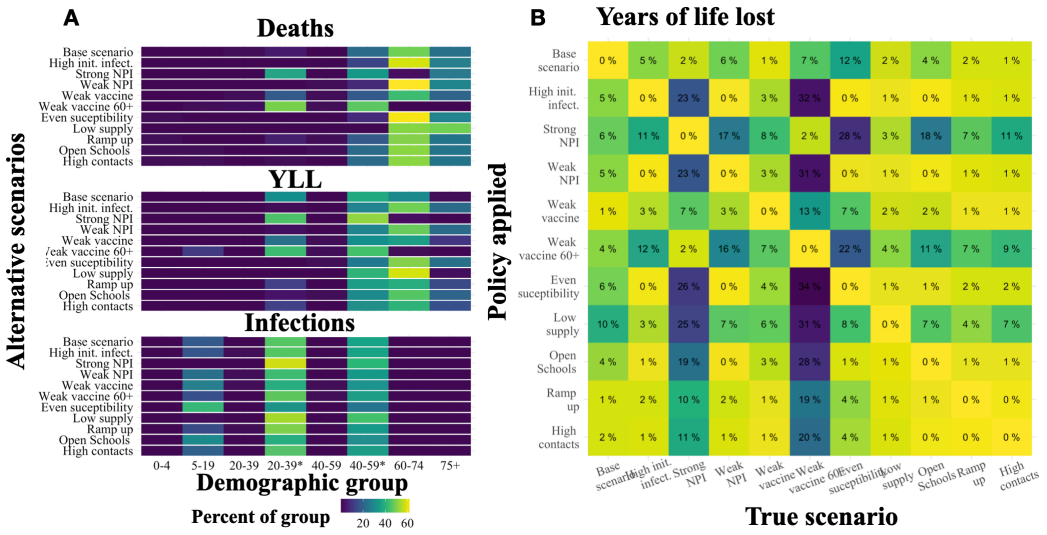

FIGURE A.2. The sensitivity of static policies to the alternative scenarios as given by the percent of the initial supply allocated to each demographic group (A), the performance relative to the optimum allocation of each policy when applied to each of the alternative scenarios when the objective is YLL (B).

A.3. Additional model robustness results

Here we present performance loss due to a mismatch between the true scenario and the one used to establish the allocation policy. Results are presented for either minimizing deaths (Fig. A.3) or infections (Fig. A.4).

A.4. Alternative model structures

In addition to considering a range of alternative parameter sets, we tested our results against three alternative model structures: (1) clustered essential workers: essential workers only contact other essential workers in the work place; (2) concentrated essential workers: 20% of the working age population are "essential" and have substantially higher (slightly above doubled) contact rates than

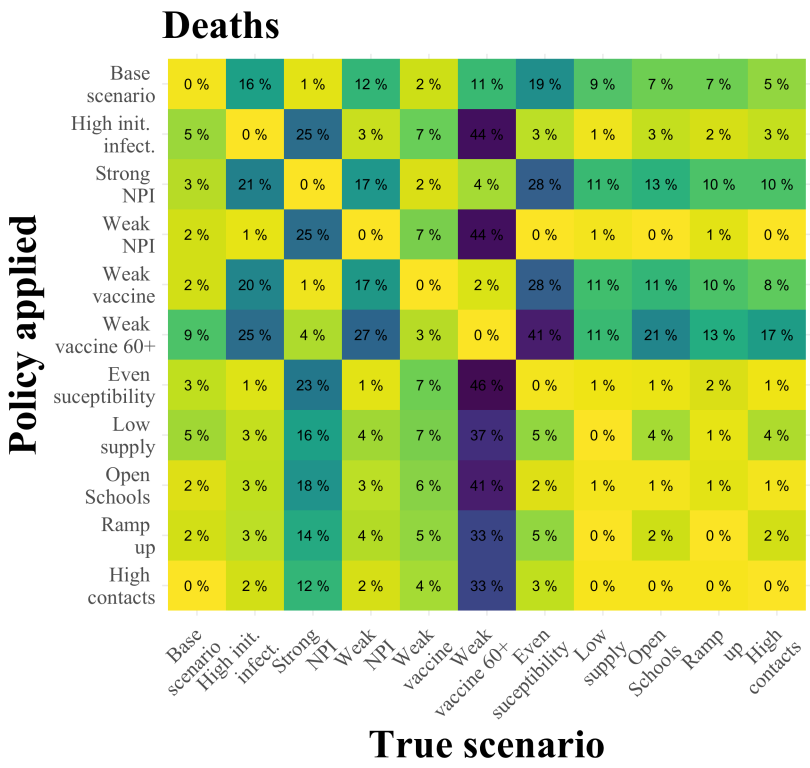

FIGURE A.3. The percentage of additional deaths in excess of the optimum when applying a policy for an alternative scenario (row) to an alternative "true" scenario (column).

essential workers in the Base model; and (3) leaky vaccine: the vaccine reduces the susceptibility of all vaccinated individuals to infection, and reduces their infectiousness and risk of death if infected. The corresponding parameter values for each alternative model are summarized below in Table A.3 and the model structure for the leaky vaccine model is given in SI Appendix A.4.1. Results presented in Fig. A.6 are discussed in the main text.

A.4.1. Leaky vaccine specification. Vaccines can provide multiple forms of protections against infections. Among these protections is the ability for vaccines to prevent individuals from becoming infected (the case considered in the main text). In addition, if vaccinated individuals still become infected they may (1) be exhibit reduced infectiousness and/or (2) develop less severe symptoms. To allow for these latter two cases we changed the model structure to track vaccinated and infected individuals. To do this we maintained the protected and uninfected category P and added four categories: vaccinated and exposed class P_{exp}, vaccinated and pre-symptomatic P_{presym}, vaccinated and asymptomatic P_{asym} and vaccinated and symptomatic P_{sym}. The effectiveness of

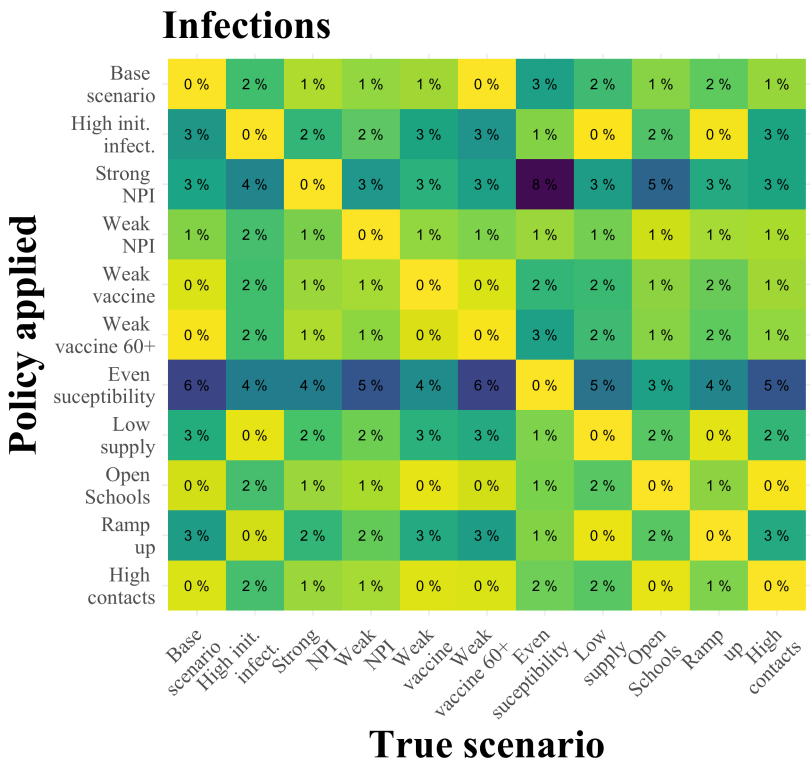

FIGURE A.4. The percentage of additional infections in excess of the optimum when applying a policy for an alternative scenario (row) to an alternative "true" scenario (column).

Model	New parameters
Clustered essential workers	Essential workers contacts focused within group as discussed in SI Appendix A.1.6
Concentrated essential workers	The proportion of essential workers is 20% work; contact rates are increased for essential workers ($\alpha_{work} = 0.1$, $\alpha_{work}* = 1.0$)
Leaky vaccine 1	$VE_{sucpt} = VE_{trans} = VE_{sym} = 0.9$
Leaky vaccine 2	$VE_{sucpt} = 0.9$, $VE_{trans} = VE_{sym} = 0$

TABLE A.3. The parameter values changed between the alternative model structures and the base model.

the vaccine is modeled with three age specific vectors, VE_{sucpt}, VE_{trans}, and VE_{sym}, which quantify the extent to which the vaccine reduces the susceptibility of vaccinated individuals to infection, the reduction in infectiousness of vaccinated individuals and the reduction in infection fatality rate of vaccinated individuals. This new model can be described by the following system of equations:

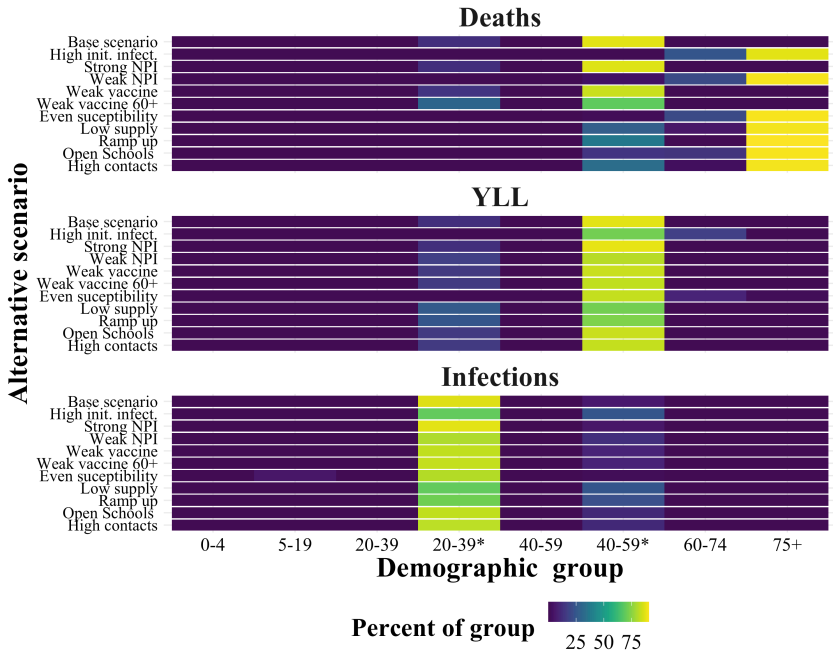

FIGURE A.5. The percentage of each demographic group vaccinated after the fist decision period under each objective and the alternative scenarios.

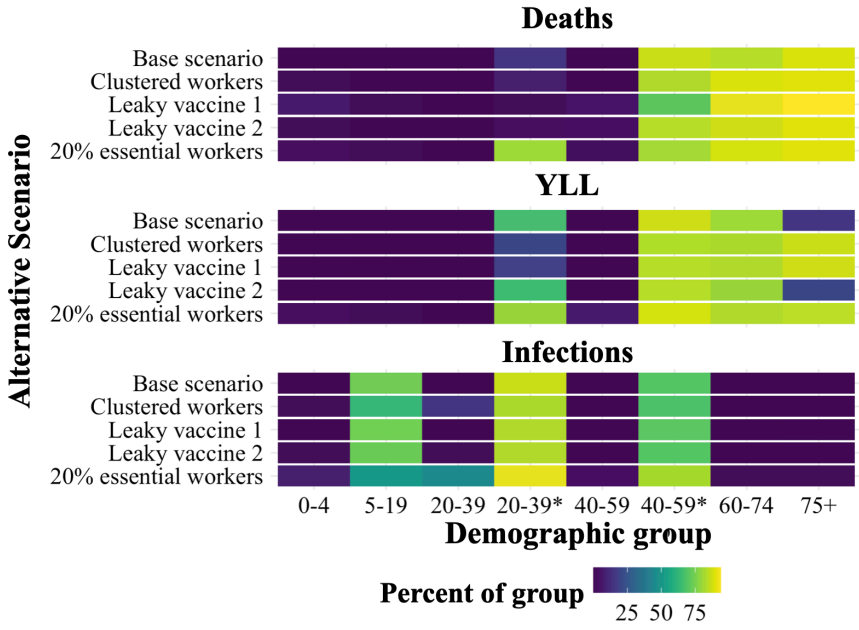

FIGURE A.6. The percent of each demographic group (horizontal axis) vaccinated after three months under the optimal policy for each of the alternative model structures (vertical axis) and objectives (panels).

$$\dot{S}_i = -qs_i\theta\left[\sum_{j\in J}\sum_{m\in M}\tau_m r_{m,i,j} S_i \frac{I_{m,j}}{N_j} + \tau_m r_{m,i,j} S_i \frac{P_{m,j}(1-VE_{trans})}{N_j}\right] - \mu_i v \tag{A.18}$$

$$\dot{P}_i = -qs_i(1-VE_{sucpt,i})\theta\left[\sum_{j\in J}\sum_{m\in M}\tau_m r_{m,i,j} P_i \frac{I_{m,j}}{N_j} + \tau_m r_{m,i,j} P_i \frac{P_{m,j}(1-VE_{trans})}{N_j}\right] + \mu_i v \tag{A.19}$$

$$\dot{E}_i = qs_i\theta\left[\sum_{j\in J}\sum_{m\in M}\tau_m r_{m,i,j} S_i \frac{I_{m,j}}{N_j} + \tau_m r_{m,i,j} S_i \frac{P_{m,j}(1-VE_{trans})}{N_j}\right] - E_i/\gamma_{exp} \tag{A.20}$$

$$\dot{P_{exp,i}} = qs_i(1-VE_{sucpt,i})\theta\left[\sum_{j\in J}\sum_{m\in M}\tau_m r_{m,i,j} P_i \frac{I_{m,j}}{N_j} + \tau_m r_{m,i,j} P_i \frac{P_{m,j}(1-VE_{trans})}{N_j}\right] - E_i/\gamma_{exp} \tag{A.21}$$

$$\dot{I}_{pre,i} = E_i/\gamma_{exp} - I_{pre,i}/\gamma_{pre} \tag{A.22}$$

$$\dot{P}_{pre,i} = P_{exp_i}/\gamma_{exp} - P_{pre,i}/\gamma_{pre} \tag{A.23}$$

$$\dot{I}_{asym,i} = \sigma_{asym} I_{pre,i}/\gamma_{pre} - I_{asym,i}/\gamma_{asym} \tag{A.24}$$

$$\dot{P}_{asym,i} = \sigma_{asym} P_{pre,i}/\gamma_{pre} - P_{asym,i}/\gamma_{asym} \tag{A.25}$$

$$\dot{I}_{sym,i} = (1-\sigma_{asym}) I_{pre,i}/\gamma_{pre} - I_{sym,i}/\gamma_{sym} \tag{A.26}$$

$$\dot{P}_{sym,i} = (1-\sigma_{asym}) P_{pre,i}/\gamma_{pre} - P_{sym,i}/\gamma_{sym} \tag{A.27}$$

$$\dot{R}_i = I_{asym,i}/\gamma_{asym} + (1-\delta_i) I_{sym,i}/\gamma_{sym} + (1-\delta_i\,(1-VE_{sym})) P_{sym,i}/\gamma_{sym} \tag{A.28}$$

$$\dot{\gamma}_i = \delta_i(1-VE_{sym}) P_{sym,i}/\gamma_{sym} + \delta_i(1-VE_{sym}) I_{sym,i}/\gamma_{sym} \tag{A.29}$$

We consider two cases: a vaccine with equal effectiveness set to 90% for consistency with the Base model $VE_{sym} = VE_{sucpt} = VE_{trans} = 90\%$, and a vaccine that only reduces susceptibility to infections $VE_{sym} = 90\%$ and $VE_{sucpt} = VE_{trans} = 0\%$.

A.4.2. Contact rates sensitivity. One key source of both uncertainty and heterogeneity between communities is the true set of underlying social contact rates. To test the effects of these parameters we considered a range of work and other contact rates around those specified in the Base scenario. The "other" (outside of the home, school and workplace) contacts were scaled from the pre-COVID-19 average while the work contact rates were increased for essential workers and held at 10% of pre-COVID-19 levels for non-essential workers. Results are presented in Fig. A.7, where each scenario is labeled with the percentage change in average work contact rate and the contact rates of essential workers are a percentage of pre-COVID-19 average levels. We found that the policies did change as both other and work contacts increased and that these changes were consistent with the changes in reproductive number induced by the higher contact rates.

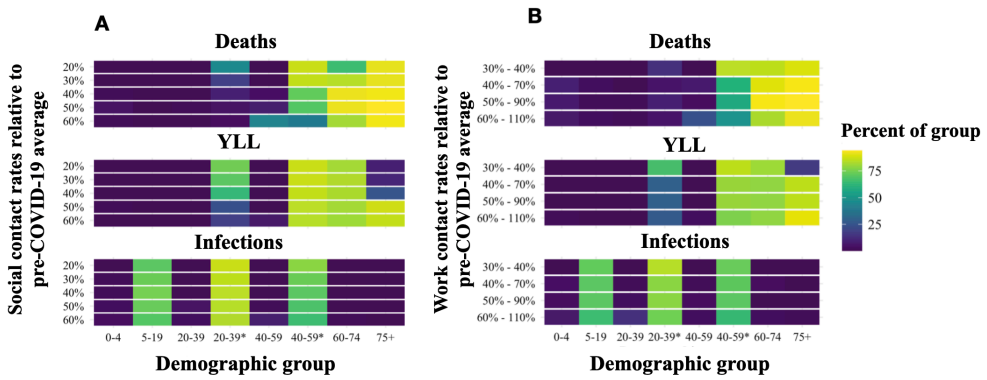

FIGURE A.7. The percentage of each demographic group vaccinated after 3 months under the optimal dynamic policy given variation in (A) "other" (non-work) contacts and (B) work contacts.

Bibliography

1. O. Byambasuren, *et al.*, Estimating the extent of true asymptomatic COVID-19 and its potential for community transmission: systematic review and meta-analysis, *Available at SSRN 3586675: https://ssrn.com/abstract=3586675* (2020).
2. A. T. Levin, K. B. Cochran, S. P. Walsh, Assessing the age specificity of infection fatality rates for COVID-19: Meta-analysis & public policy implications, *National Bureau of Economic Research Working Paper No. 27597* (2020).
3. N. G. Davies, *et al.*, *Nature Medicine* **26**, 1205 (2020).
4. S. Abrams, *et al.*, Modeling the early phase of the Belgian COVID-19 epidemic using a stochastic compartmental model and studying its implied future trajectories, *preprint*, medRxiv, doi: 10.1101/2020.06.29.20142851 (2020).
5. M. Herper, D. Garde, Moderna to submit Covid-19 vaccine to FDA as full results show 94% efficacy, STAT, https://www.statnews.com/2020/11/30/moderna-covid-19-vaccine-full-results/ (2020).
6. A. W. Bartik, Z. B. Cullen, E. L. Glaeser, M. Luca, C. T. Stanton, What jobs are being done at home during the COVID-19 crisis? Evidence from firm-level surveys, *National Bureau of Economic Research Working Paper No. 27422* (2020).
7. LMI and C2ER, SOC codes for CISA critical infrastructure workers (2020). Data retrieved from LMI institute webstite, https://www.lmiontheweb.org/more-than-half-of-u-s-workers-in-critical-occupations-in-the-fight-against-covid-19/.
8. G. McCormack, C. Avery, A. K.-L. Spitzer, A. Chandra, *JAMA* **324**, 388 (2020).
9. A. J. Kucharski, *et al.*, *The Lancet Infectious Diseases* **20**, 553 (2020).
10. H. M. Korevaar, *et al.*, Quantifying the impact of US state non-pharmaceutical interventions on COVID-19 transmission, *preprint*, medRxiv, doi: 10.1101/2020.06.30.20142877 (2020).
11. Population Pyramid, Population of the United States, https://www.populationpyramid.net, accessed July 12, 2020 (2020).
12. U.S. Social Security Administration (USSSA), Life Table (2020). Data retrieved from USSSA website, https://www.ssa.gov/oact/STATS/table4c6.html.
13. Institute for Health Metrics and Evaluation (IHME), Covid-19 mortality, infection, testing, hospital resource use, and social distancing projections (August 21, 2020 estimates), Seattle, United States of America: Institute for Health Metrics and Evaluation (IHME), University of Washington, https://ihmecovid19storage.blob.core.windows.net/latest/ihme-covid19.zip (2020).

14. B. Lovelace Jr., N. Higgins-Dunn, CDC says U.S. should have enough coronavirus vaccine to return to "'regular life"' by third quarter of 2021, cnbc.com (2020).

15. Premise, How mask-wearing has changed in America: A visual journey of increasing mask use, San Francisco, United States of America: Premise.com, https://www.premise.com/how-mask-wearing-has-changed-in-america-a-visual-journey-of-increasing-mask-use/ (2020).

16. Institute for Health Metrics and Evaluation (IHME), COVID-19: What's new for June 25, 2020, Seattle, United States of America: Institute for Health Metrics and Evaluation (IHME), University of Washington, https://www.healthdata.org/sites/default/files/files/Projects/COVID/Estimation_update_062520.pdf (2020).

17. M. Slaoui, Interview with Mary Louise Kelly: Operation Warp Speed top adviser on the status of a coronavirus vaccine, National Public Radio, https://www.npr.org/2020/09/03/909312697/operation-warp-speed-top-adviser-on-the-status-of-a-coronavirus-vaccine (2020).

18. K. Prem, A. R. Cook, M. Jit, *PLoS Computational Biology* **13**, e1005697 (2017).

19. R. Patel, I. M. Longini Jr, M. E. Halloran, *Journal of Theoretical Biology* **234**, 201 (2005).

APPENDIX B

Modeling white sturgeon hatchery impacts

B.1. Effect of parametric assumptions on simulation results

The results of our simulation analysis, and especially the link between genetic effects of the hatchery and demographic outcomes in the wild population, are determined by four biological parameters that are difficult to measure empirically: the growth rate of the wild population when rare in the absence of hatchery production R_0, the timing of natural selection with the life cycle, the curvature of the fitness function ("strength of selection") s_W, and the change in relative fitness caused by domestication selection Q_{g_1}. The effects of these parameters on the link between genetic and demographic outcomes has been explored previously (**1**, **2**), and our findings are consistent with these previous results. We illustrate their effects in this section by comparing simulations with three levels of each parameter, holding all other parameters at the base values. As expected, the quantitative predictions of the model are highly sensitive to the choice in these model parameters, but the qualitative dynamics of the fitness effects are generally consistent. In contrast, the demographic impacts of the hatchery are not always present. These are highly sensitive to the timing of selection. If all selection occurs early in the life cycle before density dependence, the effect of mortality from selection is offset by a reduction in density-dependent mortality, decoupling the fitness of the individuals from the growth rate of the population.

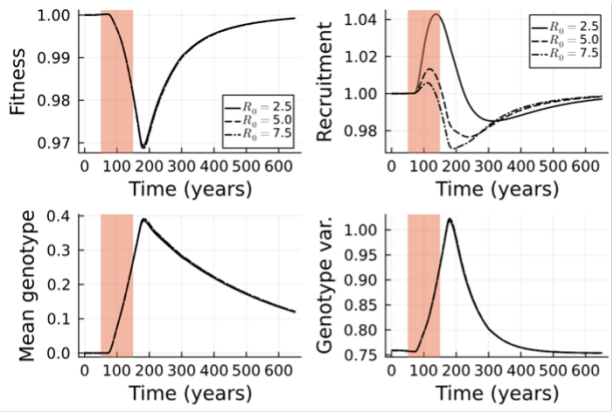

FIGURE B.1. Simulations with different values of R_0, the parameter that determines the strength of density dependence in the demographic model. Each simulation shows similar effects on fitness, but the declines in recruitment after hatchery production ends are smaller and occur after a longer delay when density dependence is weaker (Lower values of R_0).

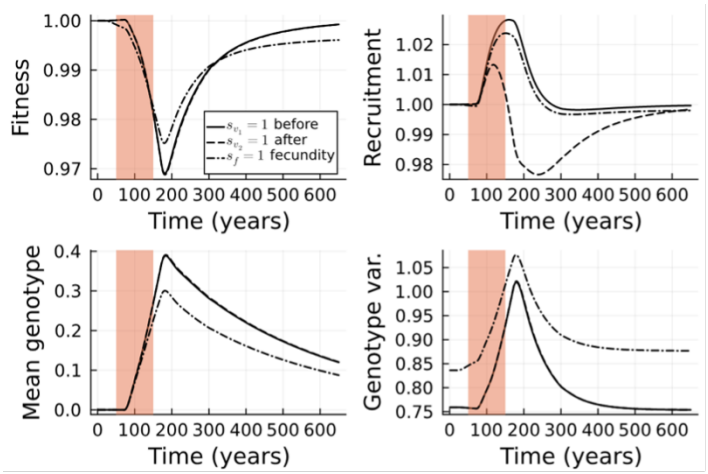

FIGURE B.2. The effect of the timing of selection on the genetic and demographic outcomes. The two viability selection events s_1 and s_2 lead to similar fitness effects, but the decline in fitness only leads to a decline in abundance if selection occurs after density-dependent mortality s_2. When selection occurs before density dependence, it reduces the level of density-dependent mortality partially offsetting the demographic impacts. Fecundity selection results in smaller fitness and demographic effects because its demographic effects are offset by density-dependent mortality and it acts directly on hatchery-origin fish in the wild reducing the rate of gene flow.

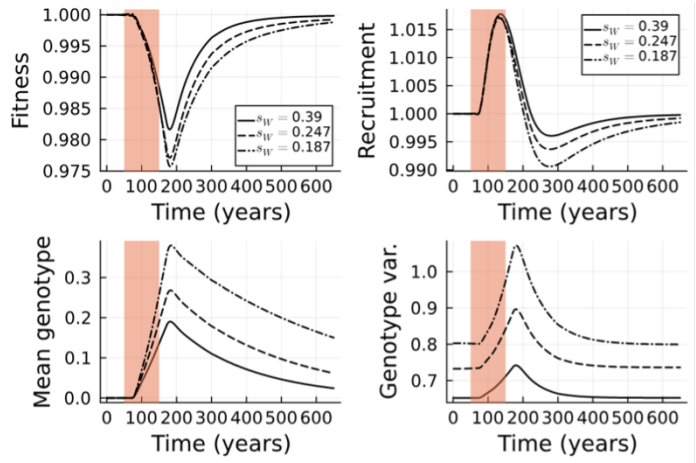

FIGURE B.3. Simulations varying the curvature of the fitness function ($s_w = s_1 + s_2 + s_f$) while holding the fitness declines in hatchery constant. More curvature (stronger selection) results in smaller declines in fitness and abundance, because stronger selection removes a greater fraction of poorly adapted individuals before they can reproduce, which in turn, reduces geneflow from the hatchery population to the wild.

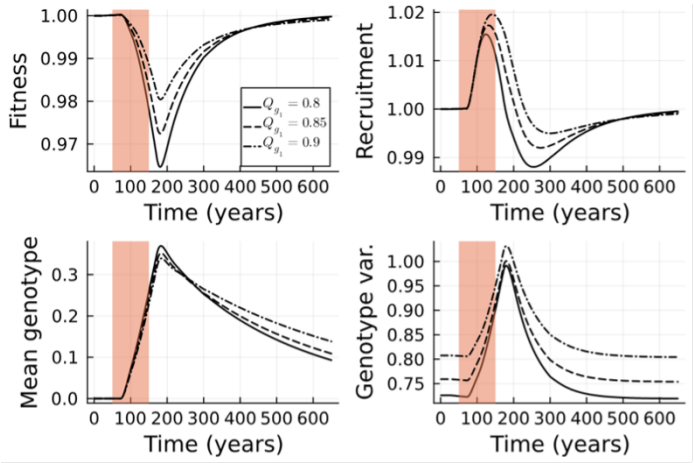

FIGURE B.4. Simulations varying the genetic component of fitness decline (values of relative fitness Q_{q_1} closer to one) in the hatchery holding the total fitness decline constant. Smaller values of relative fitness declines result in smaller impacts on fitness are recruitment in the wild population, because the process of domestication acts more slowly.

B.2. Equilibrium solution

To calibrate the model, we solved for the equilibrium abundance, age structure, and genotype distribution without immigration from the hatchery. At equilibrium, the abundance and genotype distribution of the yearling cohort will be constant over time

$$N_{W,0,t} = N^*_{W,0} \tag{B.1}$$

$$\psi_{W,0,t}(g) = \psi^*_{W,0}(g). \tag{B.2}$$

These constants then determine each age class's equilibrium abundance $N^*_{W,a}$ and genotype distribution $\psi^*_{W,a}(g)$. We calculated the equilibrium abundance and genotype distribution for each age class by tracking the changes in a cohort's abundance and genotype distribution through each step in the life cycle. The first step is density-dependent mortality (main text eqn.3.14), which affects the abundance but not the genotype distribution

$$N'^*_{W,0} = N^*_{W,0} S_0 f(\beta\ X_t) \tag{B.3}$$

$$\psi'^*_{W,0,t}(g) = \psi^*_{W,0}(g). \tag{B.4}$$

The second step is selection (main text eqn.3.10-3.11) which changes both the genotype distribution and the abundance of the cohort

$$N''^{*}_{W,0} = N'^{*}_{W,0} S^{*}) \tag{B.5}$$

$$\psi''^{*}_{W,0}(g) = \frac{1}{S^{*}} \psi'^{*}_{W,0}(g)\, W'(g) \tag{B.6}$$

$$S^{*} = \int_{-\infty}^{\infty} \psi'^{*}_{W,0,t}(g)\, W'(g)\, dg \tag{B.7}$$

We can then track the newly recruited cohort through its life history calculating the genotype distribution and abundance at each age.

$$N^{*}_{W,1} = N''^{*}_{W,0} \tag{B.8}$$

$$\psi^{*}_{W,1}(g) = \psi''^{*}_{W,0}(g) \tag{B.9}$$

$$N^{*}_{W,a+1} = S_a N^{*}_{W,a}, \quad \text{for} \quad a > 0 \tag{B.10}$$

$$\psi^{*}_{W,a+1}(g) = \psi^{*}_{W,a}(g) \tag{B.11}$$

We can then find a closed-from for the age-specific abundance and genotype distribution given the recursive definition given above in eqns.B.8-B.11

$$N^{*}_{W,a} = N^{*}_{W,0} S_0 f(\beta X_t)\, S^{*} \left(\prod_{i=1}^{a-1} S_a \right), \text{ (for } a > 1) \tag{B.12}$$

$$\psi^{*}_{W,a} is = S^{*} \psi^{*}_{W,0}(g) W'(g)\,. \tag{B.13}$$

Finally, we checked if our initial assumption that the yearling abundance and genotype distribution are constant by applying the reproduction model (main text eqn.3.1-3.5) to the equilibrium age-structured conditions defined by eqn.B.12 and B.13. This yields equilibrium conditions strictly in terms of the initial cohort abundance and genotype distribution

$$N^{*}_{W,0} = \sum_{a=1}^{A_{max}} F_a N^{*}_{W,0} S_0 f(\beta X_t)\, S^{*} \left(\prod_{i=1}^{a-1} S_a \right) \tag{B.14}$$

$$X_t = b_0 \sum_{a=1}^{A_{max}} F_a \left(\prod_{i=1}^{a-1} S_a \right) N^{*}_{W,0} S_0 f(\beta X_t)\, S^{*} \left(\prod_{i=1}^{a-1} S_a \right) + \sum_{a=1}^{A_{max}} b_a \left(\prod_{i=1}^{a-1} S_a \right) N^{*}_{W,0} S_0 f(\beta X_t)\, S^{*} \left(\prod_{i=1}^{a-1} S_a \right) \tag{B.15}$$

$$\psi^{*}_{W,0}(g) = \frac{1}{\sqrt{2\pi V_r}} \int_{\infty}^{\infty} \int_{\infty}^{\infty} e^{\frac{-(g-0.5g_1-0.5g_2)^2 2V_r \psi^{*}_{spawn}(g_1)\psi^{*}_{spawn}(g_2)}{d}} g_1 dg_2 \tag{B.16}$$

$$\psi^{*}_{spawn}(g) = \frac{1}{S^{*}} \psi^{*}_{W,0}(g) W'(g)\,. \tag{B.17}$$

To solve for the equilibrium abundance, we combine the equilibrium age structure eqn.B.12-B.13 and use the model of density dependence and selection eqn.B.3 and B.5 to rewrite the equilibrium conditions in terms of $N''^{*}_{W,0}$

$$N''^{*}_{W,0} = S^{*} S_0 f \left(\beta \left(b_0 \sum_{a=1}^{A_{max}} F_a \left(\prod_{i=1}^{a-1} S_a \right) N''^{*}_{W,0} + \sum_{a=1}^{A_{max}} b_a \left(\prod_{i=1}^{a-1} S_a \right) N''^{*}_{W,0} \right) \right) \times \sum_{a=1}^{A_{max}} F_a \left(\prod_{i=1}^{a-1} S_a \right) N''^{*}_{W,0}. \tag{B.18}$$

Solving eqn.B.18 for $N''^{*}_{W,0}$ yields the equilibrium abundance at recruitment

$$N''^{*}_{W,0} = \frac{f^{-1}\left(\frac{1}{S^{*} S_0 \sum_{a=1}^{A_{max}} F_a \left(\prod_{i=1}^{a-1} S_a \right)} \right)}{\beta b_0 \sum_{a=1}^{A_{max}} F_a \left(\prod_{i=1}^{a-1} S_a \right) + \beta \sum_{a=1}^{A_{max}} b_a \left(\prod_{i=1}^{a-1} S_a \right)}. \tag{B.19}$$

We combined eqn.B.16-B.17 to create one equilibrium conditions for the genotype distribution

$$\psi^{*}_{W,0}(g) = \frac{1}{\sqrt{2\pi V_r}} \int_{-\infty}^{\infty} \int_{\infty}^{\infty} e^{\frac{-(g-0.5g_1-0.5g_2)^2}{2V_r}} \left(\frac{1}{S^{*}} \right)^2 \psi^{*}_{W,0}(g_1) W'(g_1) \psi^{*}_{W,0}(g_2) W'(g_2) dg_1 dg_2. \tag{B.20}$$

The right-hand side of eq. S21 combines two operators, a convolution (the double integral) and multiplication by the selection gradient $W'(g)$, both of these operators preserve normal distributions (i.e. if they are applied to a normal distribution they will produce a normal distribution). Because of this property, we can check for solutions where the $\psi^{*}_{W,0}(g)$ is a normal distribution, and solve for the equilibrium mean $\hat{g}^{*}$ and variance V^{*}. The convolution operator does not change the mean of the distribution, and the selection operator pulls the mean toward the genotypic optimum therefore an equilibrium will occur at the genotypic optimum $\hat{g}^{*} = 0$. To solve for the equilibrium variance, we track the effect of the reproduction and selection operators. Selection decreases the genetic variance by a factor determined by the selection strength s. Reproduction can act to increase of decrease the genetic variance, random mating reduces the genetic variance by a factor of 1/2 because individuals with unusually large or small genotype values are more likely to mate with individuals with a less extreme genotype. However, impact heritability increases the genetic variance by a factor of V_r. Applying selection and then reproduction produces an equilibrium condition for the genetic variance of a cohort at birth

$$V^{*}_{W,0} = \frac{1}{2} \left(V^{*}_{W,0} + s \right)^{-1} + V_r. \tag{B.21}$$

solving for $V^{*}_{W,0}$ yields

$$V^{*}_{W,0} = \frac{2s}{(1 - sV_r) + \sqrt{(1 - sV_r)^2 + 4sV_r}}. \tag{B.22}$$

Therefore, an equilibrium genotype distribution is

$$\psi^{*}_{W,0}(g) = \frac{1}{\sqrt{2\pi V^{*}_{W,0}}} e^{-\frac{1}{2}\left(\frac{g^2}{V^{*}_{W,0}} \right)}. \tag{B.23}$$

B.3. Model tuning and reparameterization

B.3.1. Density dependent parameters. We tune the model to achieve a specified R_0 without altering the life history characteristics of the population by tuning the density-independent survival at age zero (S_0). The reproductive

number R_0 describes the expected number of offspring a newly born individual will have that survive to the same point in the life cycle when density-dependent competition is negligible. We calculate it by multiplying the expected lifetime egg production (LEP) of a recruit by the probability an egg survives to recruitment at the equilibrium genotype distribution

$$R_0 = S_0 S^* \times LEP = S_0 S^* \sum_{a=0}^{A_{max}} F_a \prod_{j=1}^{a} S_j. \tag{B.24}$$

Higher values of R_0 imply greater reproductive potential that is offset by density-dependent mortality at equilibrium (thus stronger larger values of R_0 imply greater density dependence). Tuning the model this way allows us to adjust the strength of density dependence without changing the equilibrium abundance, which facilitates comparisons across parameter sets. Solving for S_0 in terms of R_0 with eqn.B.24 yields

$$S_0 = \frac{R_0}{S_1^* S_2^* \sum_{a=0}^{A_{max}} F_{a,n}^* \prod_{j=1}^{a} S_j}. \tag{B.25}$$

Given a value of S_0, we solve for a the value of β that yields the target equilibrium population size using eqn.B.18

$$\beta = \frac{f^{-1}\left(\frac{S_1^*}{\sum_{a=0}^{A_{max}} F_a \prod_{j=1}^{a} S_j S_0}\right)}{N_{W,1}^* \left(b_0 S_0 \sum_{a=0}^{A_{max}} F_a \prod_{j=1}^{a} S_j + \sum_{a=1}^{A_{max}} b_a \prod_{j=1}^{a} S_j\right)}. \tag{B.26}$$

B.3.2. Trait heritability. Previous studies have shown that the trait's heritability under selection is an important driver of the genetic effects of hatchery production (**3**). However, unlike these previous studies, our model does not include heritability as a parameter. Rather it is an emergent property of the system that depends on the variance of the trait distribution $\psi_{W,a,t}$, which can vary with time. To test the effects of trait heritability on outcomes, we tune the model to fix the heritability of the trait given the equilibrium genotype distribution. The equilibrium heritability depends on the variance of environmental effects V_E and the variance of the equilibrium trait distribution V_g^*

$$h^{2^*} \equiv \frac{V_g^*}{V_E + V_g^*} \tag{B.27}$$

We derive V_g^* in Appendix 2. Given this definition, we solve for the variance of environmental effects V_E that produces the target heritability at equilibrium

$$V_E = \frac{V_g^*}{h^{2^*}} - V_g^*. \tag{B.28}$$

B.4. Numerical analysis

We used two alternative numerical techniques to represent the genotype distribution $\psi_{i,a,t}(g)$. For the simulation analysis we approximated the genotype distribution at a finite set of grid points. This method does not restrict the genotype distribution to a specific functional form and can be the most accurate representation given sufficient grid points. We also developed a method that represented the approximated the genotype distribution with a

Gaussian density function and tracked changes in the mean and variance. This approach provided a finite dimensional representation of the system that we used to calculate the Jacobian matrix for the recovery rate calculations.

B.4.1. Grid approximation. To account for deviations from normality we approximate the genotype distribution by a probability mass function over a grid of genotype values. This numerical approximation replaces the function $\psi_{i,a,t}(g)$ with a vector $\vec{\psi}_{i,a,t}$ whose entries sum to one. The selection gradients vector $\vec{W}'$ defines the fitness at each grid point. We calculated the effect of selection on the abundance, equation main text eqn. 9, with the discretized version

$$N''_{i,0,t} = N'_{i,0,t}(\vec{\psi}'_{i,a,t} \cdot \vec{W}\prime). \tag{B.29}$$

We used a similar procedure to compute the effect of selection on the genotype distribution, with a discretized version of main text equation 10

$$\vec{\psi}'_{i,a,t} = \frac{\vec{\psi}_{i,a,t} \circ \vec{W}'}{\vec{\psi}_{i,a,t} \cdot \vec{W}'}. \tag{B.30}$$

The sold dot is the usual dot product operator, and the open dot is component-wise multiplication. We calculated the abundance and genotype of the yearling age class in a similar fashion. The abundance of yearlings after reproduction is unchanged by the discretization, but the genotype distribution is. The discretized version of main text eqn.3.10 and 3.11 is

$$\vec{\psi}_{spawn,t} = \frac{1}{N_{W,0,t}} \sum_{a=0}^{A_{max}} \sum_{i \in n,h}^{F} {}_{a,o}\vec{\psi}_{i,a,t}. \tag{B.31}$$

We can calculate the genotype distribution of the yearling population in the discrete model by replacing the integrals in main text equation 3.6 with sums

$$\psi_{0,n,t}[k] \propto \sum_i \sum_j \psi_{spawn,t}[i] * \psi_{spawn,t}[j]\, e^{-\frac{(g[k]-0.5g[i]-0.5g[j])^2}{2V_r}}. \tag{B.32}$$

However, we take advantage of the structure of these sums as a convolution to increase the computational efficiency of the calculation using the Fast Fourier Transform (**4**). Eqn.B.32 is equivalent to applying two convolutions to the spawning genotype distribution $\vec{\psi}_{spawn,t}$. First the distribution of mid parental values are calculated by taking the convolution of distribution of breeding values in the spawning population $\vec{\psi}_{b,t}$ with itself where breeding values are exactly one half an individuals genotype value $\vec{\psi}_{b,t}[x] = \vec{\psi}_{spawn,t}[2x]$. Next, we then apply a convolution with the inheritance kernel $\vec{K}$, a discretized normal distribution with mean zero and variance V_r, to the result of the first operator to obtain the final value of the genotype distribution.

$$\psi_{0,n,t} = (\vec{\psi}_{b,t} * \vec{\psi}_{b,t}) * \vec{K}, \tag{B.33}$$

where $*$ denotes the convolution operator.

B.4.2. Normal approximation. Another way to represent the genotype distribution numerically is to approximate it with a normal distribution, where we then follow the expectation and variance of the genotype distribution. In the following section we compute the effects of selection, reproduction and gene flow on the expectation and variance of the genotype distribution, assuming the trait distribution normal.

B.4.2.1. *Selection.* When the genotype distribution $\psi'_{W,0,t}$ is normal with mean $E_{W,0,t}$ and variance $V_{W,0,t}$ the equation for selection (main text 3.10, 3.11) are

$$N'_{W,0,t} = \int_{-\infty}^{\infty} \frac{1}{sV_E+1} \sqrt{\frac{1}{2\pi V_{W,0,t}}} e^{-\frac{1}{2}\left(\frac{s\left(g-p_{opt}\right)^2}{2(sV_E+1)} + \frac{\left(g-E_{W,0,t}\right)^2}{2V_{W,0,t}}\right)} dg \tag{B.34}$$

$$\psi'_{W,o,t}(g) = \frac{N'_{W,0,t}}{N_{W,0,t}} \frac{1}{sV_E+1} \sqrt{\frac{1}{2\pi V_{W,0,t}}} e^{-\frac{1}{2}\left(\frac{s\left(g-p_{opt}\right)^2}{2(sV_E+1)} + \frac{\left(g-E_{W,0,t}\right)^2}{2V_{W,0,t}}\right)}. \tag{B.35}$$

To compute the new abundance and the expectation and variance after selection we rearranged the exponents in eqns.B.34 and B.35 to take the form

$$\frac{(g-\mu)^2}{2\sigma^2} + c. \tag{B.36}$$

Rewriting the equations in this way reveals that they are equivalent to a normal distribution times a constant. The expectation of the updated distribution is μ, the variance is σ^2 and the constant c determines the survival rate. Using this approach, we found that the updated expectation and variance of the genotype distribution are given by

$$E'_{W,0,t} = \frac{V_{W,0,t}}{V\prime_{W,0,t}} E_{W,0,t} \tag{B.37}$$

$$V'_{W,0,t} = \frac{V_{W,0,t}^{(sV_E+1)}}{sV_{W,0,t}^{+sV_E+1}}. \tag{B.38}$$

The abundance after selection is determined by the abundance before selection and the constant c applying this yields the updated abundance

$$N'_{W,0,t} = N_{W,0,t} \sqrt{\frac{V'_g}{V_g}} e^{-\frac{1}{2}\left(\frac{E^2}{V_g}\left(1-\frac{V'_g}{V_g}\right)\right)}. \tag{B.39}$$

B.4.2.2. *Reproduction.* In place of main text equations 3.6 and 3.7, which calculate the effects of reproduction on the entire genotype distribution, here we calculate the mean and the variance of the genotype distribution of the yearling age class. The mean is given by the weighted sum of the mean genotypes in each of the spawning age classes.

$$E_{W,0,t} = \sum_{a=0}^{A_{max}} F_a N_{W,a,t} + \sum_{a=0}^{A_{max}} F_a N_{H,a,t}. \tag{B.40}$$

We compute the variance of the yearling age class in two steps. First we computed the variance of the mid parental genotype values of spawning pairs V_b, because we assume that spawning pairs form at random this value is half the

variance of the genotype distribution in the spawning population

$$V_{b,t} = \frac{1}{2}\left(\sum_{a=1}^{A_{max}} \sum_{o\in\{W,H\}} F_a {V'}_{o,a,t} + \sum_{a=1}^{A_{max}} \sum_{o\in\{n,c\}} F_a {E_{o,a,t}}^2 - \left(E_{W,0,t}\right)^2\right) \tag{B.41}$$

To account for the variability in the inheritance of genotype values we then add the recombination variance to this value

$$V_{W,0,t} = V_{b,t} \ + V_r. \tag{B.42}$$

B.5. Marginal effect of removals

The effect of removing a hatchery-origin individual on the genetic state of the population depends on its age. For example, removing a young (immature) individual is likely to have a smaller effect than removing an older mature individual because they are more likely to die before they have a chance to reproduce. We can formalize this idea by calculating the expected lifetime reproductive success of an individual from the life table parameters for an individual of age a

$$E_a\left[F\right] = \sum_{i=a}^{A_{max}} \left(F_i \times \prod_{j=a}^{i} S_j\right). \tag{B.43}$$

Where F_i is the fecundity at age i and $\prod_{j=a}^{i} S_j$ is the probability of surviving from age a to age i.

Bibliography

1. F. Débarre, S. Gandon, *The American Naturalist* **177**, E84 (2011).
2. M. L. Baskett, R. S. Waples, *Conservation Biology* **27**, 83 (2013).
3. M. J. Ford, *Conservation Biology* **16**, 815 (2002).
4. M. Turelli, N. H. Barton, *Genetics* **138**, 913 (1994).

APPENDIX C

Fisheries Monitoring

C.1. Sensitivity Analysis

C.1.1. Minimum Biomass. One important assumption of the model is the minimum biomass. I assume that the population cannot be driven to extinction by the fishery, which is consistent with both theory and empirical evidence (**1**). However, this assumption requires a fixed lower bound on the population size B_{min} In the main text I set this value at $5\%B_{MSY}$. To test the sensitivity of the results to this assumption I calculated the marginal value of monitoring in the base scenario for three values $\{0.01B_{MSY}, 0.025B_{MSY}, 0.05B_{MSY}\}$. The results of this analysis are shown in SI figure C.1. The effect of B_{min} has a small influence on the value of monitoring over most of the state space, but significantly increases the value of monitoring for high values of uncertainty CV_t.

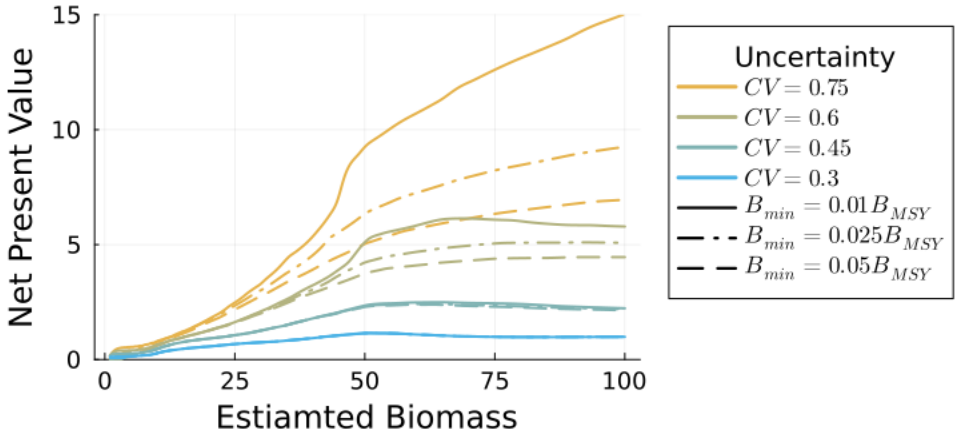

FIGURE C.1. SI Figure 1.1) The effect of B_{min} on the value of monitoring as a function of the estimated biomass (x-axis) and uncertainty (colors). The values of B_{min} are denoted dashed and solid lines.

C.1.2. Incorporating uncertainty in the harvest control rule. In our base specification, the uncertainty buffers are fixed and do not scale with uncertainty CV_t. However, some fisheries do account for the level of uncertainty in their harvest control rules. One simple way to incorporate uncertainty is the $p^* - \sigma$ rule. This rule is defined by asking a simple question: Given our uncertainty about the abundance of the stock "σ," what is the probability of overfishing given a harvest limit $\overline{H}_t$? The harvest limit is then chosen to limit the probability of overfishing below a target value p^*.

Overfishing occurs when the realized exploitation ratio $F_t = \overline{H}_t/B_t$ exceeds the target exploitation ratio $\overline{F}_t$. I define the target exploitation ratio using an HCR similar to the one defined in the base specification (eqn. 4.6), but the

value depends on the true biomass of the stock B_t

$$\overline{F}_t(B_t) = \begin{cases} F_{target} & \text{if } B_t > B_{target} \\ Z_t F_{target} + (1 - Z_t) F_{dm} & \text{if } B_{limit} < B_t \leq B_{target} \\ F_{dm} & \text{if } B_t \leq B_{limit} \end{cases} \tag{C.1}$$

$$Z_t \equiv \frac{B_t - B_{limit}}{B_{target} - B_{limit}}. \tag{C.2}$$

Notice that both the realized exploitation ratio F_t and the target exploitation ratio depend on (unobserved) biomass of the stock B_t. As a result, these values are not known with certainty. Instead, we can model them as random variables whose distributions depend on the manager's beliefs about the biomass of the stock $\hat{B}_t$ and CV_t. To find the harvest limit we need to solve for the harvest limit $\overline{H}_t$ where the probability F_t exceeds $\overline{F}_t$ is equal to p^*

$$p^* = P(F_t > \overline{F}_t | \hat{B}_t, CV_t, \overline{H}_t). \tag{C.3}$$

In the numerical implementation, I solve equation C.3 for $\overline{H}_t$ using the bisection algorithm.

To show how this specification changes the results, I computed the value of monitoring information and the monitoring frequency sensitivity analysis. The marginal benefit and the total value of monitoring are both lower when the $p^* - \sigma$ rule is used(fig.C.2). The effect is greater for the total value. This difference can likely be explained by the difference in performance in the absence of monitoring. Int he base scenario, the harvest control rule is set based on the expected abundance of the stock even when uncertainty is very high, resulting in a high probability of over fishing and destroying the value of the stock. in contrast the $p^* - \sigma$ rule systematically adds a buffer, reducing the probability of over fishing while still maintaining some harvest in the long run. The largest effect of the $p^* - \sigma$ rule was reducing the sensitivity of the optimal policy to the non-consumptive values. This is likely because the highest probability of depleting the stock occurs when uncertainty is high. However, when the $p^* - \sigma$ rule is used, the harvest limit is significantly reduced when uncertainty is high, mitigating the risk of depleting the stock.

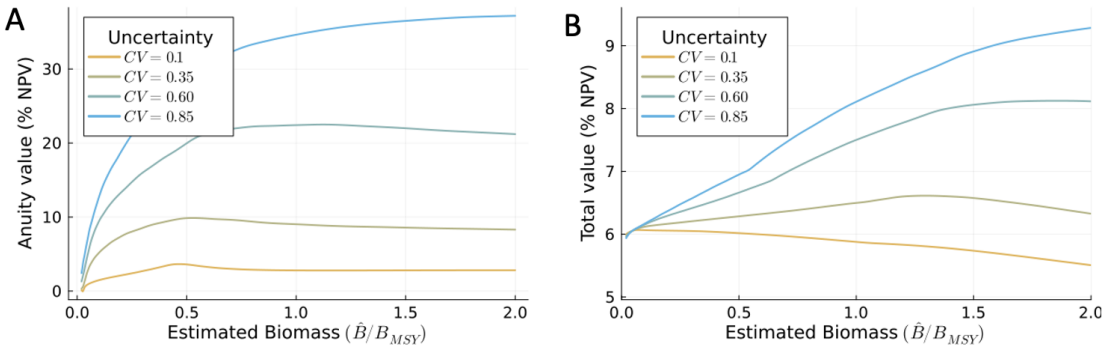

FIGURE C.2. A) Importance of each model parameter for determining the frequency of monitoring under the optimal policy. B) The partial dependence of the frequency of monitoring on each model parameter.

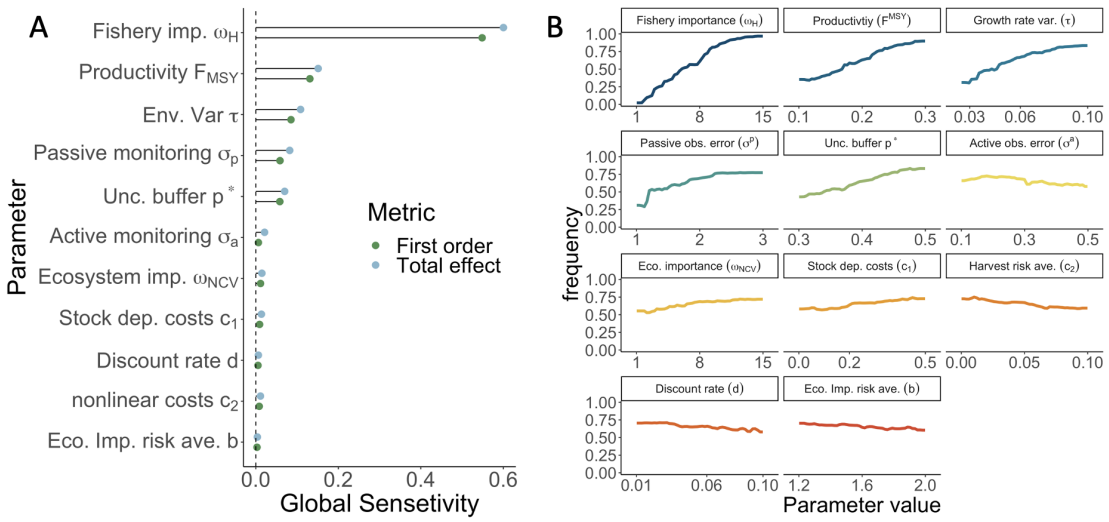

FIGURE C.3. A) Importance of each model parameter for determining the frequency of monitoring under the optimal policy. B) The partial dependence of the frequency of monitoring on each model parameter.

C.2. Approximation of the manager's problem

The manager's problem is a partially observable Markov decision process (**2**). We approximate the solution by restating the problem as a belief-state Markov decision process (MDP), which we solve using value function iteration (VFI). Standard MDP solution techniques are limited when the observations y_t are noisy because observations from previous time steps $y_{t-\tau}$ contain information about the system's current state B_t not captured by the most recent observation y_t. Belief state MDPs solve this problem by defining the optimal policy as a function from the manager's beliefs about the system's state $f(B_t)$ to an action i_t. However, solving for the exact optimal policy is infeasible because the belief states f_t is a probability distribution, an infinite dimensional mathematical object. Therefore, we choose to approximate the solution to the manager's problem by representing the manager's belief state f with a parametric approximation; we then use the parameters of this distribution to define the optimal policy (**3**). Following (**4**, **5**), I approximate the belief state with a log-normal density function with parameters μ_t and η_t, I find the best fitting parameters to the true belief state f using density projection (**3**) which, in this, case correspond to the expectation $\mu_t = E_f[log(B_t)]$ and variance $\eta_t^2 = V_f[log(B_t]$ of the log biomass. To improve interpretability, We present the value and policy functions in terms of the expected biomass $E[B] = \hat{B} = e^{\mu+0.5\eta^2}$ and coefficient of variation $CV = \sqrt{e^{2\eta} - 1}$.

I solve for the optimal policy using the Bellman equation, which implicitly defines the expected net present value of the fishery as a function of the belief state variables $V(\mu_t, \eta_t^2)$. The bellman equation defines the value function and the expected return from the fishery in the current period $E[\pi(B_t, H_t, i_t)|\mu_t, \eta_t^2]$ plus the discounted expected value

of the fishery in the following year $E[V(\hat{B}_{t+1}, CV_{t+1}|\mu_t, \eta_t^2]$

$$V(\hat{B}_t, CV_t) = \max_{i_t \in \{0,1\}} \left\{ E[\pi(H_t, B_t, i_t)|\mu_t, \eta_t^2] + \delta E[V(\mu_{t+1}, \eta_{t+1}^2)|\mu_t, \eta_t^2, i_t] \right\}. \tag{C.4}$$

The optimal policy is defined as the argument i_t that satisfies the maximization operator in the Bellman equation

$$P(\hat{B}_t, CV_t) = \underset{i_t \in \{0,1\}}{argmax} \left\{ E[\pi(H_t, B_t, i_t)|\mu_t, \eta_t^2] + \delta E[V(\mu_{t+1}, \eta_{t+1}^2)|\mu_t, \eta_t^2, i_t] \right\}. \tag{C.5}$$

The integrals used to compute the expectations operators in equations C.4 and C.5 are defined in SI appendix C.2.1. In short, these operators integrate over the belief state distribution $f(B_t)$ to account for uncertainty in the current biomass and integrate over the probability of transitioning between pairs of belief states $T(\mu_{t+1}, \eta_{t+1}^2|\mu_t, \eta_t^2, i_t)$ to account for uncertainty in the future state of the system. I solve equation C.4 using the value function iteration algorithm, approximating $V(\mu_t, \eta_t^2)$ at a fixed set of regularly spaced grid points defined by pairs (μ, η^2). I approximated values between the grid points by interpolating them with a cubic b-spline. Details of this implementation are given in SI appendix C.2.2, and code for the analysis is available on GitHub.

C.2.1. Belief state transition probabilities. As discussed above, I approximate the belief state $f_t(B_t)$ with a log-normal distribution and track the changes in the mean μ_t and variance η_t^2 of the log biomass x_t. The mean and variance are updated in two steps corresponding to equations 4.5 and 4.4. The first step incorporates information from new observations using Baye's law, and the second accounts for the effect of harvest and population growth. The manager's prior beliefs $f(B_t)$ and the measurement errors are normally distributed (after a log transformation). Normal priors and likelihoods form a conjugate family of distributions; therefore, the Bayes' update changes the parameters of the prior distribution but not the functional form. We apply these formulas to the initial mean and variance μ_t, η_t^2 to obtain the updated values

$$\mu_t{}' = \left(\frac{\mu_t}{{\sigma_{i_t}}^2} + \frac{y_t}{\eta_t^2} \right) \left(\frac{1}{{\sigma_{i_t}}^2} + \frac{1}{\eta_t^2} \right)^{-1}, \tag{C.6}$$

$$\eta_t^{2\prime} = \left(\frac{1}{{\sigma_{i_t}}^2} + \frac{1}{\eta_t^2} \right)^{-1}. \tag{C.7}$$

The second step, which calculates the effect of harvest and population growth on abundance, does not have a closed-form solution because the stock dynamics are non-linear. Instead, the updated mean μ and variance η^2 are represented as integrals that calculate the expectation of the biomass after the nonlinear growth equations have been applied

$$\mu_{t+1} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \phi(x_t|\mu_t', {\eta_t'}^2) \phi(\nu_t|0, \tau^2) log \left(\frac{r\left(e^{x_t} - H_t\right) e^{\nu_t - 0.5\tau^2}}{1 - b\left(e^{x_t} - H_t\right)} \right) dx_t d\nu_t, \tag{C.8}$$

$$\eta_{t+1}^2{}' = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \phi(x_t|\mu_t', {\eta_t'}^2) \phi(\nu_t|0, \tau^2) \left(log \left(\frac{r\left(e^{x_t} - H_t\right) e^{\nu_t - 0.5\tau^2}}{1 - b\left(e^{x_t} - H_t\right)} \right) - \mu_{t+1}' \right)^2 dx_t d\nu_t. \tag{C.9}$$

where $\phi(x|\mu, \eta^2)$ is the normal density function with mean μ and variance η. In the following sections, having a compressed notation for each of these operations will be convenient. In the following, I refer to the Baye's update

equations (eqn. C.6) as a function $U_b(\mu, \eta^2, y, i)$, and the population growth update (eqn. C.8 as the function $U_g(\mu, \eta^2)$.

C.2.2. Value function iteration. The Value Function Iteration algorithm has two primary components: a method for calculating the expectation operators in the Bellman equation (eqn. C.4) and a method to numerically approximate the value function $V(\mu_t, \eta_t^2)$

C.2.3. Expectation operators. The Bellman equation C.4 evaluates the expected reward in the current period, given the uncertainty in the biomass, and the expected value of the fishery in the following period, given the uncertainty in the biomass and growth rate. The first term is computed by integrating the objective function $\pi(H_t, B_t, i_t)$ over the normal approximation of the managers belief state

$$E[\pi_t|\mu_t, \eta_t^2] = \int_{-\infty}^{\infty} \pi(H_t, e^{x_t}, i_t)\phi(x_t|\mu_t, \eta_t^2)dx_t \tag{C.10}$$

I approximate equation C.10 with Gauss-Hermite weights ω_i and nodes n_i. Gauss Hermite quadrature integrates functions of the form $f(x_t)e^{-x^2}$. Therefore, I used a change of variables to calculate the expectation of $\pi(H_t, B_t, i_t)$ with respect to the belief state distribution

$$E[\pi_t|\mu_t, \eta_t^2] = \frac{1}{\eta_t\sqrt{2\pi}} \int_{-\infty}^{\infty} \pi(H_t, e^{x_t}, i_t)e^{-\frac{1}{2}\left(\frac{x_t-\mu_t}{\eta_t}\right)^2} dx_t \tag{C.11}$$

$$u = \eta_t x_t + \mu_t \tag{C.12}$$

$$E[\pi_t|\mu_t, \eta_t^2] = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \pi(H_t, e^{\frac{u-\mu_t}{\eta_t}}, i_t)e^{-\frac{1}{2}u^2} du \tag{C.13}$$

$$\approx \frac{1}{\sqrt{2\pi}} \sum_{i=0}^{N} \pi(H_t, e^{\frac{n_i-\mu_t}{\eta_t}}, i_t)e^{\frac{1}{2}n_i^2}\omega_i \tag{C.14}$$

The second expectation in equation C.4 accounts for uncertainty in the values of the belief state variables μ_{t+1}, η_{t+1}^2 in the following period. The belief state transitions are defined in SI appendix C.2.1 as a function $U(\mu, \eta^2, y_{t+1}, i_t)$. The only source of uncertainty in the belief state transition is the observation y_{t+1}. To compute the expected value of the belief state after the updating step, we need to describe the probability distribution of y_{t+1} given the current belief state variables μ_t, η_t^2. I calculate it in two steps. First, I compute the prior distribution for biomass at the time when the observation is made $f_t'(B_{t+1})$ by applying the growth update $U_g(\mu, \eta^2)$ (eqn.C.8) to obtain the mean μ_t' and variance $\eta_t^{2\prime}$. Second, I combine the prior distribution $f_t'(B_{t+1})$ with the observation $g(y_t|B_t)$ (eqn.4.3) model to get the prior predictive distribution

$$h_t(y_{t+1}) = \int_0^{\infty} f_t'(B_{t+1})g(y_{t+1}|B_{t+1})dB_t \tag{C.15}$$

$$h_t(y_{t+1}) = \int_0^{\infty} \phi(x_{t+1}|\mu_t', \eta_t^{2\prime})\phi(y_{t+1}|x_{t+1}, {\eta_{i_t}}^2)dx_t \tag{C.16}$$

$$h_t(y_{t+1}) = \phi(y_{t+1}|\mu_{t+1}', {\eta_{t+1}^2}' + {\sigma_{i_t}}^2). \tag{C.17}$$

Given the prior predictive distribution the expected value in the following period if found by integrating over the possible values of y_t

$$E[V(\mu_{t+1}, \eta^2_{t+1})] = \int_{-\infty}^{\infty} V(U(\mu_t, \eta^2_t, y_t, i_t))\phi(y_t|\mu'_t, \eta^{2'}_t + {}_{i_t}{}^2)dy_t \tag{C.18}$$

$$E[V(\mu_{t+1}, \eta^2_{t+1})] \approx \frac{1}{\sqrt{2\pi}} \sum_{i=0}^{N} V\left(U\left(\mu_t, \eta^2_t, \frac{n_i - \mu'_{t+1}}{\eta'_{t+1}}, i_t\right)\right) e^{\frac{1}{2}n_i^2}\omega_i. \tag{C.19}$$

C.2.4. Value function approximation. I solved the bellman operator with an integrative procedure called value function iteration (VFI). The algorithm initialized the value function V at a fixed set of grid points in the belief state space, in this case, pairs of mean and variance values (μ_t, η^2_t). I chose a set of grid points evenly spaced with $n = 50$ grid point in each dimension resulting in a total of $n^2 = 2,500$ points. Grid point in the mu_t dimension spanned from $0.05B_{MSY}$ to $2B_{MSY}$ and from 0 to $log(2)$ in for σ^2_t, which corresponds to a coefficient of variation of one. I interpolated the values between the grid points using cubic b-splines.

The value function was initialized at zero at each grid point. The VFI algorithm then updated the value at each grid point by evaluating the bellman operator. This procedure repeats until the changes to the value function fall below a threshold value of 10^-4.

C.3. Testing the solution approximation

C.3.1. Monte Carlo simulations. I used simulations to test the performance of the monitoring policies, compute the frequency of monitoring under the optimal policy, and visualize the dynamics of the fishery. I initialized each simulation by sampling the true biomass B_0 from the initial belief state distribution $f(B_0)$. I simulated the biomass of the stock forward in time by applying equations 4.1 and 4.2. The observations y_t are drawn from the observation model (eqn. 4.3) given the true biomass B_t.

The manager's belief state was simulated using the particle filter algorithm. This approach differs from the density projection technique I used in the optimization algorithm. I used this method for the simulation because it approximates exact Bayesian inference. However, it was to computationally expensive to use in the optimization procedure. This algorithm is initialized by sampling N_{pf} particles (i.e. possible values of B_0) from the initial belief state $f(B_0)$. Two operations are applied to these particles each period, corresponding to equations 4.5 and 4.4, respectively. First, each particle is updated with equations 4.1 and 4.2 to capture the effects of equation 4.5 on the belief state. Next, each particle is weighted by the likelihood of the observation y_t (eqn. 4.3). The particles are then re-sampled with replacement according to these weights. The estimated biomass is defined as the sample mean of the particles and

the coefficient of variation is computed from the sample variance

$$\hat{B}_{pf,t} = \frac{1}{N_{pf}} \sum_{i=0}^{N_{pf}} B_{t,i} \tag{C.20}$$

$$CV_{pf,t} = \frac{\sqrt{\frac{1}{N_{pf}} \sum_{i=0}^{N_{pf}} \left(B_{t,i} - \hat{B}_{pf,t} \right)^2}}{\hat{B}_t} \tag{C.21}$$

(C.22)

The harvest limit $\overline{H}_t$ and monitoring decision i_t are found by applying the harvest control rule and monitoring policy to the mean $\hat{B}_{pf,t}$ and coefficient of variation $CV_{pf,t}$ of the particle distribution. The realized harvest H_t is found by applying the harvest constraints to the true biomass B_t and harvest limit $\overline{H}_t$.

I computed the rewards for each period R_t, by applying equations 4.9 and 4.10 to the true biomass B_t and realized harvest H_t. I computed the realized present value for each simulation by summing the discounted rewards $\delta^t R_t$. I repeated this procedure N_{MC} times to estimate the expected net present value

$$ENPV = \frac{1}{N_{MC}} \sum_{i=0}^{N_{MC}} \sum_{t=0}^{T} \delta^t R_t \tag{C.23}$$

C.3.2. performance tests. To test the value function iteration algorithm and density projection approximation of the belief state I compared the value function produced by the VFI algorithm to the expected net present value of the system computed by Monte Carlo simulation using the base parameter set. I computed the expected net present value at nine points in the belief state space using three values for the expected biomass $\hat{B}_t \in \{15, 50, 115\}$ and three values for the coefficient of variation $CV_t \in \{0.25, 0.5, 0.75\}$. The value function and simulation-based estimates showed close agreement C.4. However, the value function was systematically lower than the simulation-based estimates. This difference in value may be caused by differences in the methods used to approximate the belief state transitions (eqn. 4.5 and 4.4). The VFI algorithm assumes that the belief state is always normally distributed, an approximation that is relaxed by the particle filtering algorithm used in the simulation. The management strategy may perform slightly better in the simulations because the harvest level and monitoring decisions are based on a more accurate representation of the belief state.

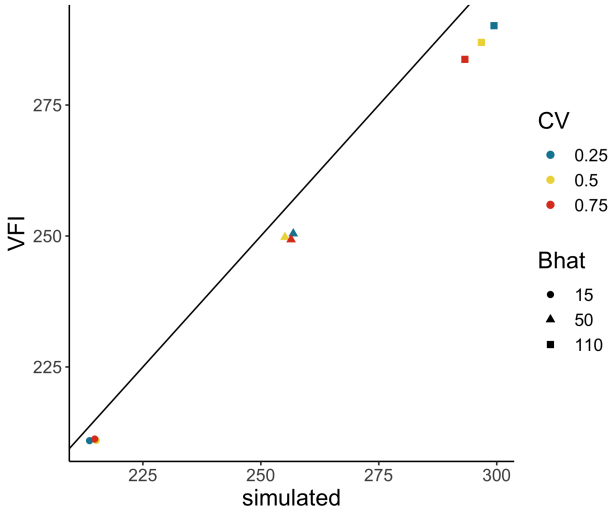

FIGURE C.4. Comparison of the net present value under the optimal policy computed by value function iteration using the density projection approach to approximate the belief state dynamics and the net present value computed by Monte Carlo simulation using a particle filtering algorithm to represent the beleif state.

C.4. Q function

The Q function defines the expected net present value of the fishery conditional on the manager's choice in monitoring decision i_t. This value can be calculated by removing the maximization operator from the Bellman equation

$$Q(\hat{B}, CV_t, i_t) = E\left[\pi(H_t, B_t, i_t)\middle|\hat{B}_t, CV_t\right] + \delta E\left[V(\hat{B}_{t+1}, CV_{t+1})\middle|\hat{B}_t, CV_t\right]. \tag{C.24}$$

C.5. Supplemental figures

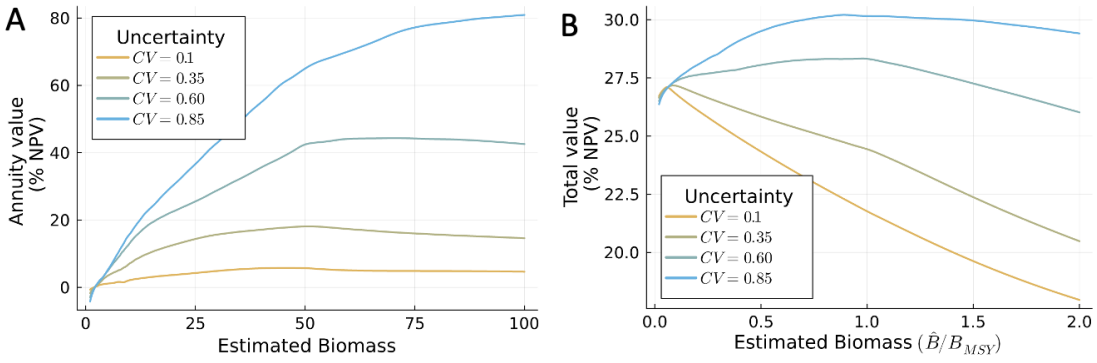

FIGURE C.5. The economic value of the monitoring program in the base scenario. A) The marginal benefit of monitoring as a percentage of the net present value of the fishery, and B) the total value of the monitoring program as a percentage of the net present value.

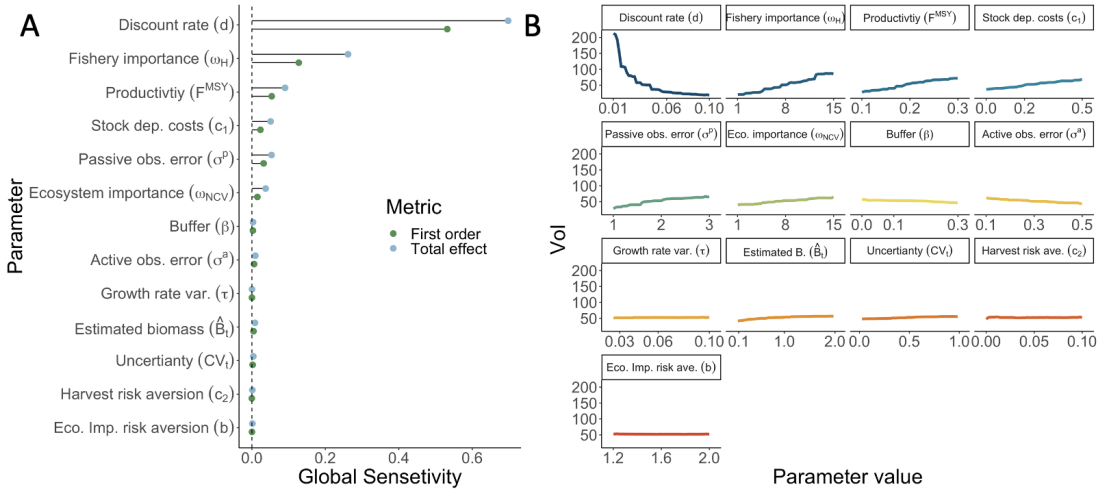

FIGURE C.6. A) Importance of each model parameter and the belief state variables for determining the total value of the monitoring program. B) The partial dependence of the value of monitoring on each belief state variable and model parameters.

Bibliography

1. O. Le Pape, S. Bonhommeau, A.-E. Nieblas, J.-M. Fromentin, *Proceedings of the National Academy of Sciences* **114** (2017).
2. P. L. Fackler, R. G. Haight, *Resource and Energy Economics* **37**, 226 (2014).
3. E. Zhou, M. Fu, S. Marcus, *Automatic Control, IEEE Transactions on* **55**, 1101 (2010).
4. M. R. Sloggy, D. M. Kling, A. J. Plantinga, *Journal of Environmental Economics and Management* **103**, 102357 (2020).
5. D. M. Kling, J. N. Sanchirico, P. L. Fackler, *Journal of Environmental Economics and Management* **84**, 223 (2017).