Title
IMPLICATIONS OF POLICY ADJUSTMENT COSTS FOR FISHERIES MANAGEMENT

Journal
Natural Resource Modeling, 30(1)

ISSN
0890-8575

Authors
RYAN, DANIEL
TOEWS, CARL
SANCHIRICO, JAMES N
et al.

Publication Date
2017-02-01

DOI
10.1111/nrm.12107

Peer reviewed

NATURAL RESOURCE MODELING
Volume 00, Number 0, July 2016

IMPLICATIONS OF POLICY ADJUSTMENT COSTS FOR FISHERIES MANAGEMENT

DANIEL RYAN
National Institute for Mathematical and Biological Synthesis, Knoxville, TN 37996
E-mail: danryanfl@yahoo.com

CARL TOEWS
Department of Mathematics and Computer Science, University of Puget Sound, Tacoma, WA 98416
E-mail: ctoews@pugetsound.edu

JAMES N. SANCHIRICO
Department of Environmental Science and Policy, University of California, Davis, CA 95616
E-mail: jsanchirico@ucdavis.edu

PAUL R. ARMSWORTH*
National Institute for Mathematical and Biological Synthesis, Knoxville, TN 37996 and Department of Ecology and Evolutionary Biology, University of Tennessee, Knoxville, TN 37996
E-mail: p.armsworth@utk.edu

ABSTRACT. Optimal policy recommendations from natural resource management models can involve fast-changing management interventions. Actual policies often change more gradually, potentially reflecting costs associated with policy adjustment. We examine how including policy adjustment costs changes policy recommendations in models of fishery management. Specifically, we examine cases where policy adjustment costs increase at an increasing rate with the rate of change of fishing effort levels. We examine how accounting for these costs changes optimal management strategies in static and time-varying environments. Increasing policy adjustment costs results in optimal approach paths to stationary optimal solutions that are more gradual but that rely on oscillatory convergence. In time-varying environments, including policy adjustment costs smoothes optimal management recommendations, resulting in decreased variation in effort levels on the stationary part of the optimal solution, but increased variation in optimal stock sizes. We conclude that accounting for policy adjustment costs in this way leads to management recommendations that in some ways would be simpler for managers to implement (more gradual approach paths), but in other ways more challenging (oscillatory convergence) and that can affect where environmental variability is expressed in model predictions (stock sizes versus effort levels) while not suppressing this variability altogether.

KEY WORDS: Regularization, harvesting, institutions, inflexible, sluggish, recruitment variation, implementation.

1. Introduction. One of the most important insights in the management of natural resources was the framing of the problem as an investment decision for society (Clark and Munro [1975], see also Clark [1990] and for a recent review, Conrad

*Corresponding author: Paul Armsworth, Department of Ecology and Evolutionary Biology, 569 Dabney Hall, 1416 Circle Drive, University of Tennessee, Knoxville, TN 37996, USA. e-mail: p.armsworth@utk.edu.

Received by the editors on 6th October 2015. Accepted 29th June 2016.

and Smith [2012]). The management strategies suggested by bioeconomic models that formalize this framing often entail making rapid changes to harvesting policies. For example, bioeconomic models often suggest that making fast transitions from some initial ecosystem state toward a long-term management target would be optimal (e.g., Spence and Starrett [1975], Hartl and Feichtinger [1987], Sanchirico et al. [2010], Martinet et al. [2010]). Also, when managing stochastically varying populations, the classic optimal management strategy is that harvest rates in each period track inter-annual variability in recruitment closely (Reed [1979], Nostbakken and Conrad [2007], McGough et al. [2009], Kapaun and Quaas [2013]). In another example, some deterministic problems, for example, how to manage age-structured populations with harvesting strategies that cannot target individual cohorts, involve recommendations that pulse harvesting would be optimal (Pope [1973], Bjorndal and Brasao [2006], Tahvonen [2009], Armsworth et al. [2011]). With a pulse harvesting strategy, a population is harvested intensely for a short period and then left unharvested for a number of years to recover before being harvested intensely once more. As well as policies that change quickly in time, policies that change frequently in space have also been recommended based on optimal control applications (e.g., Neubert [2003], Neubert and Herrera [2008]).

Fishery management policies, such as changes to total allowable catches, however, rarely follow such fast changing policy prescriptions (Biais [1995], Patterson and Resimont [2007], Boettiger et al. [2016]). Rather, policies are sticky or slow to change suggesting the presence of adjustment or transaction costs associated with rapid changes in optimal policies. There are many political-economy reasons for the presence of these costs. For example, there might be social pressures to maintain fleet sizes when population sizes are low due to concerns about social and economic displacement (Clark et al. [1985]). The latter is evident in many fishery management debates surrounding rebuilding of fish populations (National Resource Council [2014]) and is the reason for disaster relief packages being provided to fishing industry participants. Other reasons could be associated with the time scale of stock assessments (e.g., required every 3 years) or fishery management plans (e.g., setting policies for a number of years), or with uncertainty surrounding assessments of current cohort sizes (Ralston et al. [2011]).

In this paper, we set out to investigate how including adjustment costs changed the predictions of a canonical deterministic optimal control model of fishery management (see Boettiger et al. [2016] for an investigation of these costs in a stochastic fishery model). Specifically, we incorporate a penalty into the objective function that increases in the change in fishing effort (denoted $(dE/dt)^2$ penalty). Inclusion of penalty functions of this type, termed *regularization*, is common in applications of control theory in engineering (Kirsch [1996]). Additional costs associated with quickly changing harvesting effort levels would emerge from increased administrative transaction costs associated with repeatedly revisiting past policy decisions, or from preferences on the part of managers, fishermen, fish processing plants, or other stakeholders for more regular harvests and employment schedules. A few studies have utilized this type of representation of policy adjustment costs usually in the course of examining other fisheries management questions (Ludwig [1980], Wirl [1992], Jørgensen and Kort [1993], Feichtinger et al. [1994], Liski et al. [2001]). Our goal in this work is to focus on the consequences of policy adjustment costs *per se* to aid natural resource modelers in understanding the consequences of including them in a model.

The paper is structured as follows: the "Methods" section records our modeling and parameter choices, and briefly describes our numerical methods. We focus on deterministic optimal control via harvesting in a constant environment and in one where stock productivity is characterized by periodic variation as a consequence of oscillating environmental conditions (e.g., El Nino-La Nina cycles; see also Parma [1990], Castilho and Srinivasu [2007], Carson et al. [2009]). The "Results" section presents and analyzes solutions in both environmental settings. These results have an added complexity in the time-varying setting, in which a central question is how the time scale of the control interacts with the speed of the periodic variation. The "Discussion" section summarizes management implications and provides recommendations for the design of natural resource management models. The Supporting Information contains additional details relevant to our analysis, as well as further numerical results that complement and refine results in the main text.

2. Methods.

The following notation will be used throughout the paper:

Symbol	Meaning
t	time
$N(t)$	biomass of fish stock at time t
$E(t)$	total effort invested in fishing at time t
$h(t)$	harvest rate
r or $r(t)$	intrinsic rate of increase of the fish stock (see below)
K	carrying capacity (constant)
p	price of fish (constant)
q	catchability coefficient (constant)
δ	discount rate (constant)

Other symbols will be defined as they are introduced.

2.1. Growth and harvest rates.

Throughout the paper, the harvest rate is represented by a Schaefer model, that is, is proportional to both stock size and fishing effort via the equation

$$h(t) = qE(t)N(t).$$

Stock growth is assumed to be logistic, with intrinsic rate of increase r and carrying capacity K,

$$\frac{dN(t)}{dt} = rN(t)\left(1 - \frac{N(t)}{K}\right) - qE(t)N(t). \tag{1}$$

As noted in the notation table, carrying capacity K and catchability coefficient q are assumed constant. We start by assuming the same for growth parameter r. Then, we examine how different policy adjustment costs affect the ability of the optimal solutions to track changing environmental conditions by considering a case where $r = r(t)$ varies sinusoidally with time,

$$\frac{dN(t)}{dt} = r(t)N(t)\left(1 - \frac{N(t)}{K}\right) - qE(t)N(t), \tag{2}$$

where $r(t)$ is of the form

$$r = r(t) \equiv 1 - 0.5\cos\left(\frac{\pi t}{L}\right)$$

for some constant L. Note that with this formulation, the population always has a positive intrinsic rate of increase, but environmental conditions oscillate between favorable conditions during which the population is highly reproductive and less favorable conditions where its reproductive potential is lower.

2.2. Management objectives. As a baseline, we assume the net present value (NPV) of future harvests is given by

$$NPV_1 \equiv \int_0^T e^{-\delta t}\left[pqE(t)N(t) - c_1 E(t)\right]\, dt\,, \tag{3}$$

where T is a time horizon and c_1 is a constant that determines the cost of effort. Here the managerial objective is to determine the value of $E(t)$ that maximizes (3), subject to stock dynamics (1). It is well known (see, e.g., Spence and Starrett [1975] and Clark [1990]) that the linearity of the objective in the control leads to a "bang-bang" solution, that is, one in which stock levels follow a most rapid approach path (MRAP) to and from the steady-state solution, and the effort function exhibits discrete jumps. We have chosen to show results for a finite time problem to highlight more clearly how including policy adjustment costs affects approach paths to stationary optimal solutions. We utilize time horizons long enough that the relevant intuition for the infinite time horizon problem can be gained by ignoring the terminal time approach path.

To account for policy adjustment costs, we assume that these costs scale quadratically with the rate of change of effort, where the NPV function is

$$NPV_2 \equiv \int_0^T e^{-\delta t} \left(pqE(t)N(t) - c_1 E(t) - c_2 \left(\frac{dE(t)}{dt} \right)^2 \right) dt \,, \tag{4}$$

where $c_2 > 0$ is a constant. In this formulation, the state becomes a two-dimensional vector $(N(t), E(t))$, and the control is the function $dE(t)/dt$. This formulation assumes that the policy adjustment costs are symmetric with respect to increases or decreases of effort, which is a reasonable starting place as fishermen often argue that fishing harvests are slow to increase and conservationists argue that fishing catches are slow to decrease. Considering asymmetric policy adjustment costs (e.g., where it is more difficult to reduce effort than to increase it) would provide a sensible extension.

Quadratic scaling is not essential, but it does represent a simple, analytically tractable way to impose the extra costs, and thus a reasonable point of departure for this study. The two different management objectives (3) and (4), when constrained by the two versions of the stock dynamics (1) and (2), describe four different maximization problems. For each, we require that $E(t) \geq 0$ and $N(t) \geq 0$. In all cases, we assume that the initial value of the stock is given by $N_0 = K$.

2.3. Parameter choice. We focus on qualitative changes in the optimal management recommendations and their impact on state dynamics under each of the different objective-state equation pairs. We emphasize in particular how optimal management recommendations are affected as the importance of additional costs in determining the overall management objective is increased, as determined by parameter c_2 in NPV_2. To produce figures illustrating the types of qualitative changes that we observe as c_2 is varied, other parameter values were chosen to be illustrative only. Specifically, unless otherwise stated, we set $r = 1$, $K = 100$, $q = 1$, and $p = 1$. We used a discount rate $\delta = 0.05$ and set the linear cost of effort $c_1 = 15$ such that under open access conditions, the fishery would be harvested to 15% of virgin biomass with linear profits.

2.4. Solution techniques. The solution to NPV_1 with state dynamics (1) is given in (Clark [1990]), and a solution with state dynamics (2) is derived in the Supporting Information.

The solutions with NPV_2 are complicated by the fact that the objective is not linear in the control. We employ a *direct method* to find these solutions, that is, we discretize time to turn the control problem into a finite dimensional nonlinear optimization problem, which we then solve with an sequential quadratic programming (SQP) method (see Buskens and Maurer [2000] for a nice overview of this approach). Specifically, we choose $N+1$ uniformly spaced points t_i on $[0,T]$, and seek an optimal control among the set of functions that are continuous and linear between these t_i. Since each such control is uniquely defined by its values at the t_i, this restriction reduces the problem to a finite-dimensional one amenable to the techniques of nonlinear programming. Starting with a discrete control vector, we compute the objective functional by integrating the state equations forward in time using a fourth-order Runge Kutta scheme. The gradient of the objective function is then calculated numerically, and used as the basis for an update to the control vector, a process that repeats until convergence is achieved. The numerical solutions displayed in this paper were achieved by formalizing this discrete optimization problem and feeding it into Matlab's `fmincon` function, a standard nonlinear constrained optimization solver.

A strength of the direct approach is that it uses only the control variables as the optimization variables and dispenses completely with the need to calculate the adjoint. (The adjoint can be regained after the algorithm has converged, however; see Buskens and Maurer [2000] for details.) One of its weakness, however, is that it can be analytically opaque. Since we wish to study the equilibrium solutions, we also employ an *indirect method* in which we introduce an adjoint variable and invoke Pontryagin's Maximum Principle to convert the problem into a two-point boundary value problems (see, e.g., Lenhart and Workman [2007].) This transformation allows us to use local linearization analysis to analyze long-run equilibria.

3. Results.

This section presents graphical, numerical, and analytical comparisons of solutions to the two different management objectives (3) and(4) under the two models for stock dynamics (1) and (2).

3.1. Static growth parameter.

When stock dynamics are logistic, that is, given by (1), optimal solutions to each management objective (3) and (4) involve three distinct phases. First, the optimal stock and effort are characterized by an initial transition from the starting condition of the system toward a long-term target for management. Then there is a period of management about this long-term target (henceforth the stationary part of the optimal solution). Finally, there is another transition from the long-term target toward a population size at which net revenue from harvesting would be zero (often termed open access or bionomic equilibrium) at the terminal time. (This latter transition would be absent in an infinite time horizon problem.)

In our baseline problem, we can be more specific, because NPV_1 is linear in the control, $E(t)$. For linear control problems, the initial and terminal transitions are

MRAPs. For instance, for a previously unharvested population, the optimal policy entails a pulse of intense fishing effort that instantaneously moves the stock from its initial, unharvested state to the desired stationary stock size. Another pulse of intense fishing effort is also applied at the terminal time to move the stock size instantaneously to the biomass level corresponding to open access conditions where instantaneous profit is zero, $N(T) = c_1/pq$. In between these two pulses of effort, the stationary part of the optimal policy is equilibrial with a constant population size, harvest, and effort level through time. See Clark [1990] for details, proofs, and additional figures.

The standard approach for solving a control problem of the form (4) is to cast it as a two-state optimal control problem. Here, we set one state to be the stock level N and the other state effort level E, and define the control to be $u = \frac{dE}{dt}$. With this convention, the optimal control problem becomes

$$\begin{aligned} &\max_u \int_0^T e^{-\delta t}\left(pqEN - c_1 E - c_2 u^2\right) dt \text{ subject to} \\ &\dot{N} = rN\left(1 - \tfrac{N}{K}\right) - qEN\,, \quad \dot{E} = u\,, \\ &N(0) = N_0\,, \quad N, E \geq 0. \end{aligned} \tag{5}$$

The current value Hamiltonian[1] in this case is

$$H^c = pqEN - c_1 E - c_2 u^2 + \lambda_1 \left[rN\left(1 - \frac{N}{K}\right) - qEN \right] + \lambda_2 u\,,$$

where λ_1 and λ_2 satisfy

$$\begin{aligned} \dot{\lambda_1} &= \delta\lambda_1 - \left[pqE + \lambda_1 r - \frac{2r\lambda_1 N}{K} - \lambda_1 qE \right], \\ \dot{\lambda_2} &= \delta\lambda_2 - [pqN - c_1 - \lambda_1 qN]\,, \\ &\lambda_1(T) = 0\,, \quad \lambda_2(0) = \lambda_2(T) = 0\,. \end{aligned}$$

The Maximum Principle implies that if N^*, E^*, λ_1^*, and λ_2^* are the optimal states and adjoints, then the optimal control u^* must satisfy

$$\frac{\partial H^c}{\partial u}(N^*, E^*, \lambda_1^*, \lambda_2^*, u^*) = -2c_2 u^* + \lambda_2^* = 0\,,$$

whereupon solving for u^* yields

$$u^*(\lambda_2^*) = \frac{\lambda_2^*}{2c_2}\,.$$

As before, we can now eliminate the control from the state and adjoint system to arrive at the optimality system

$$\dot{N} = rN\left(1 - \frac{N}{K}\right) - qEN\,, \tag{6}$$

$$\dot{E} = \frac{\lambda_2}{2c_2}\,, \tag{7}$$

$$\dot{\lambda_1} = \delta\lambda_1 - pqE - \lambda_1 r + \frac{2r\lambda_1 N}{K} + qE\lambda_1\,, \tag{8}$$

$$\dot{\lambda_2} = \delta\lambda_2 - pqN + c_1 + \lambda_1 qN\,, \tag{9}$$

$$N(0) = N_0\,, \quad \lambda_1(T) = 0\,, \quad \lambda_2(0) = \lambda_2(T) = 0\,.$$

We solve for the steady-state solution by setting the left-hand side of this system equal to zero and solving for the four unknowns N, E, λ_1, and λ_2. From (7) it is clear that $\lambda_2^* = 0$, and using this fact, together with some algebra, we find that N^* is the positive root of the quadratic polynomial

$$N^{*2} \cdot 2qrp + N^* \cdot (pKq(\delta - r) - rc_1) - \delta c_1 K. \tag{10}$$

The quantities E^* and λ_1^* are then given in terms of N^* by

$$E^* = \frac{r}{q}\left(1 - \frac{N^*}{K}\right) \tag{11}$$

and

$$\lambda_1^* = p - \frac{c_1}{N^*},$$

respectively.

Figures 1(a) and (b) illustrate full optimal solutions to objective (4) for three different values of c_2. In line with this derivation, the steady-state stock size (10) and effort levels do not depend on c_2. This feature is a consequence of the fact that a constant growth rate implies a constant harvest rate at equilibrium, at which point the penalty term dE/dt disappears entirely.

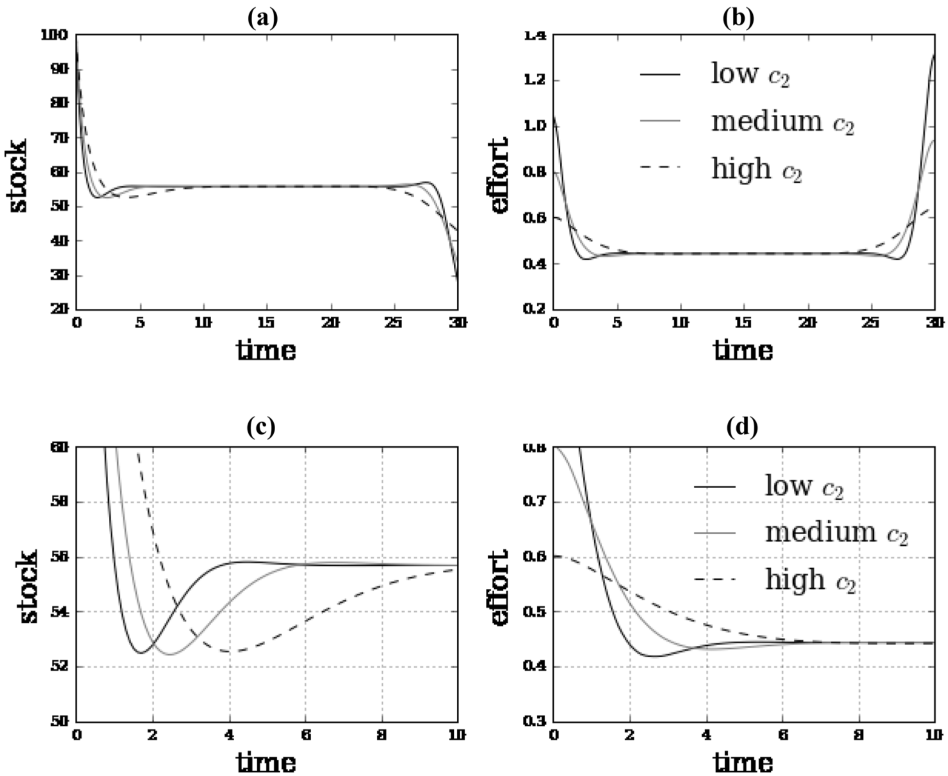

FIGURE 1. Stock sizes (panels a and c) and effort levels (panels b and d) through time corresponding to optimal management in the presence of policy adjustment costs associated with the rate of change of fishing effort. Black, gray and dashed curves correspond to increasing cost coefficients, c_2, in each case. The solutions shown in panels a and b are shown over a shorter time period in panels c and d to illustrate approach paths more clearly. State dynamics given by logistic growth with a constant intrinsic rate of increase r. Parameters: $[r, K, q, p, c_1, \delta] = [1, 100, 1, 1, 15, 0.05]$.

On the other hand, c_2 *does* influence the approach path, which as can be seen in Figures 1(c) and (d) is oscillatory. In other words, the optimal policy would have managers first overshoot the long-term target by maintaining a higher effort level for longer, before compensating by reducing the effort level below the long-run target resulting in an oscillatory approach. Mathematical justification for this oscillation can be achieved by the local linearization analysis of the associated boundary value problem. The Jacobian for the system (6)–(9) is a four by four matrix in the seven problem parameters r, p, q, K, δ, c_1, and c_2. Algebraic expressions for the eigenvalues are daunting, so we solve for the eigenvalues numerically as a function of c_2, setting all the other parameters to their usual values, that is, $p = r = q = 1$, $K = 100$, $\delta = 0.05$, and $c_1 = 15$. For each value of c_2, we find one complex

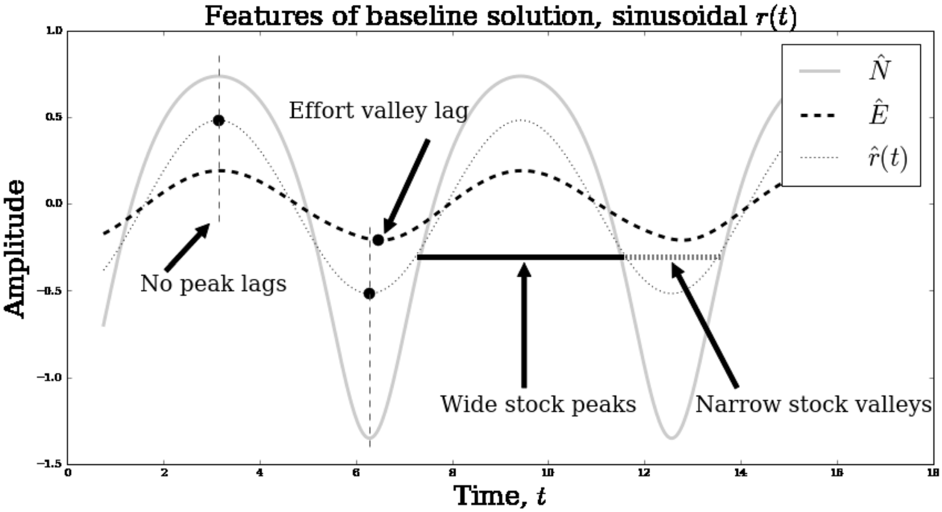

FIGURE 2. A snapshot of the singular optimal stock and effort solution, normalized to have zero mean and superimposed over $r(t)$, for the case of time-varying growth ($r(t)$). The annotations draw attention to alignment (or nonalignment) between peaks and valleys, and the phenomenon of wide peaks in the stock signal. Parameters: $[r, K, q, p, c_1, \delta] = [1, 100, 1, 1, 15, 0.05]$.

conjugate pair with negative real part, and one such pair with positive real part. The imaginary parts of the complex conjugate pairs describe the oscillatory approach while the real parts describe the rate of approach.

3.2. Time-varying growth parameter. When stock dynamics vary with time, the periodic features of the growth signal are reflected in the solution signals. The "baseline" problem for this analysis is linear management objective (3) with stock dynamics (2). An analytic solution to this problem is worked out in the Supporting Information and is similar to the solution for the case of constant r but includes some more complex functional forms. Nonetheless, because the objective is linear in E, the solution will be again an "MRAP" path, that is, the stock will descend rapidly from N_0 until it meets the singular path, and then follow the singular path until it departs, where the departure too will be as rapid as possible. Figure 2 shows a view of the singular stock and effort solutions superimposed over a plot of $r(t)$. In this plot, all signals have been translated to have mean zero, so it is easy to compare amplitude and peak or valley positions. Note that the stock and effort functions are oscillatory, with the same period as $r(t)$, but that the stock function is asymmetric, with wide peaks and narrow valleys. It turns out this particular asymmetry diminishes as c_1 increases (Supporting Information). Note too that effort and stock peaks coincide almost exactly with those of $r(t)$, but that

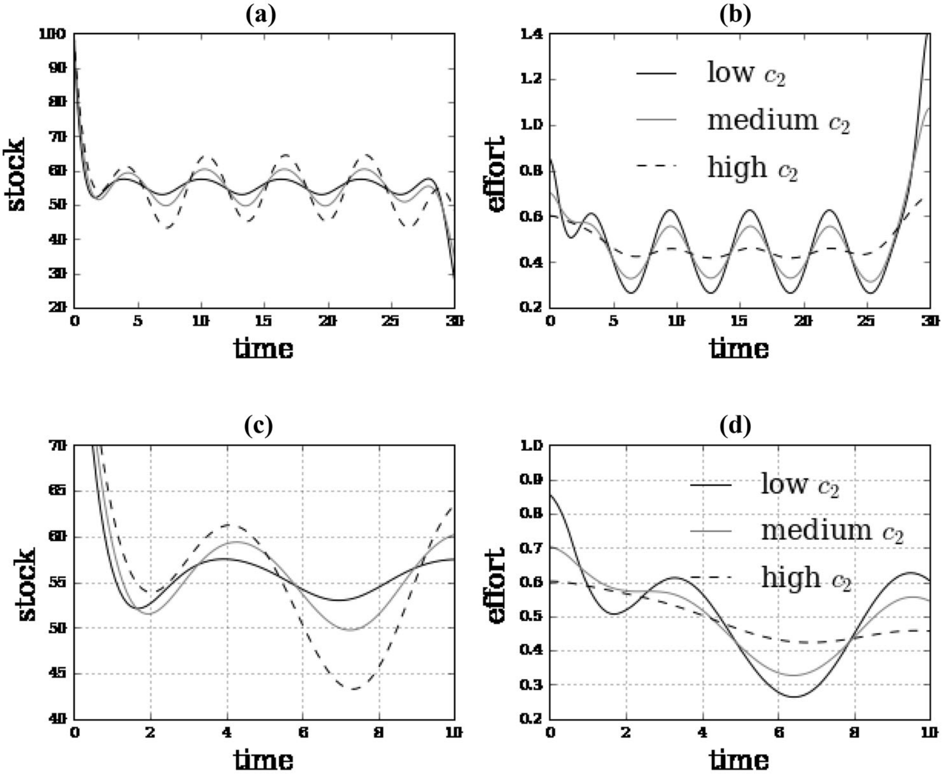

FIGURE 3. Stock sizes (panels a and c) and effort levels (panels b and d) through time that correspond to optimal management in the presence of policy adjustment costs associated with the rate of change of fishing effort. Progression from solid black, to gray to dashed curves illustrates how optimal stock sizes and effort levels change as the severity of policy adjustment costs increases. State dynamics are given by logistic growth with a varying intrinsic rate of increase $r(t)$; for this illustration, $r(t)$ is described using a sinusoid with a period of six time steps. The solutions shown in panels a and b are shown over a shorter time period in panels c and d to illustrate approach paths more clearly. The average value and amplitude of r are, respectively, set at 1 and 0.5. Parameters: $[K, q, p, c_1, \delta] = [100, 1, 1, 15, 0.05]$.

effort valleys are slightly out of phase. Once again, this phase shift depends on c_1, as shown in Figure S2.

Under the regularization scheme (4), the problem becomes nonlinear in the control, and there are changes in the approach paths and stationary solution. Some of these changes closely mirror those of the static r case. Figure 3 shows full solution trajectories for both stock and effort under variable growth model (2) and a broad range of regularization parameters c_2. Comparing Figure 3 to Figure 1, we see that in both cases the regularization scheme succeeds in smoothing out transitions to

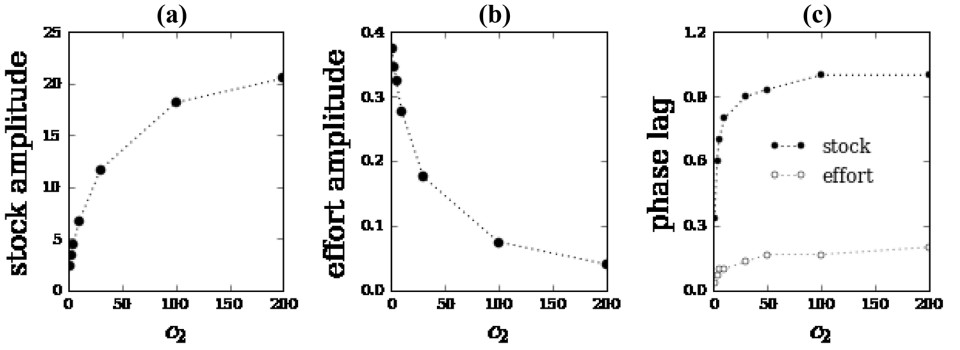

FIGURE 4. Amplitude and phase metrics for optimal singular paths in a time-varying environment in the presence of policy adjustment costs associated with the rate of change of fishing effort. Amplitude was calculated as the average difference between local maxima and minima. Peak lag was calculated as the average difference between the position of peaks in $r(t)$ and the nearest peaks in the solution signals. Note that as the regularization constant grows, stock variation rises, effort variation shrinks, and peak lags grow. For further analysis of the periodic properties of the solution signals, see the Supporting Information, especially Figures S1 and S2. Parameters: $[K, q, p, c_1, \delta] = [100, 1, 1, 15, 0.05]$.

and from the stationary path. The "overshoot" phenomenon observed for (4) is still present but in Figure 3 it is obscured by signal oscillations.

A more significant distinction between Figures 1 and 3 is that the latter contains amplitude and phase features, with a period that exactly matches that of the growth model. Figure 3 shows that as c_2 increases, the amplitude of the stock variation increases and the amplitude of the effort variation decreases. A close inspection of the figure also reveals that changes in the regularization parameters cause small shifts in the alignment between peaks and valleys. Figure 4 summarizes these results by plotting the magnitude of the amplitude and phase shift against the size of the regularization parameter, showing, for example, that very large values of c_2 have the effect of suppressing the variation in effort levels almost completely. Phase shifts, which represent the temporal lags between the growth signal and the solution signals, increase monotonically with the regularization parameter. The Supporting Information contains further discussion of signal shape, alternative phase measures, and convergence in solution-space as the regularization terms goes to zero.

4. Discussion.

Policy adjustment costs are likely to be common in many natural resource management settings. Indeed, in our own experience, recommendations that would require managers to implement fast changing policies (e.g., bang-bang approach paths or pulse harvesting) are often questioned on the grounds that such policies would be difficult, that is, expensive, to implement. Here, we examined one way that policy adjustment costs could be included in harvesting models: by attaching additional costs to quick changes to harvesting effort.

As one would expect, including policy adjustment costs results in a more gradual approach path to the stationary solution from the initial or terminal conditions being optimal. Also, accounting for policy adjustment cost leads to a smoothing of effort levels through time along the stationary part of the path in time-varying environments. In many settings, these two changes to optimal policy recommendations are consistent with observations of resource managers' actions.

However, if a modeler associates these costs with the rate of change of effort motivated due to a desire to "simplify" a policy recommendation (e.g., by smoothing effort levels or avoiding a most rapid approach recommendation), other forms of policy complexity can arise. In particular, when policy adjustment costs are associated with $(dE/dt)^2$, the optimal approach path is oscillatory, something that could also be very difficult to implement.

Moreover, a necessary counter-part to the inclusion of policy adjustment costs smoothing optimal effort levels in time-varying environments is that the stock sizes will be more variable in time. In essence, the control variables absorb less of the recruitment variability inherent in the stock. Having more variable stock sizes is less of a concern in simple deterministic models of the type that we consider, but may be important when moving to stochastic settings, especially if there are Allee thresholds, extinction boundaries, or other undesirable system states to be avoided (Roughgarden and Armsworth [2001], Boettiger et al. [2015]).

While we assumed policy adjustment costs increase at an increasing rate with the rate of change of fishing effort levels, other ways of representing these costs also make sense. For example, we compared our results with a formulation in which the marginal cost per unit of effort increases in effort by including a quadratic penalty on effort itself E^2. This alternative approach is more commonly motivated not from a policy adjustment perspective but rather from underlying characteristics of the fishing catch methods and operations that can lead to congestion externalities (e.g., Brown [1974], Lewis [1981], Anderson [1982], Hanson and Ryan [1998], Friedland [2010], Kellner et al. [2011]). Optimal policies with a penalty on E^2 share many characteristics with those that we found when the penalty is attached to $(dE/dt)^2$. Approach paths were again more gradual, but this time were monotonic and not oscillatory. Also, the variation in effort levels decreased and that in the stock size increased with the severity of the penalty on effort in the time-varying case. In one important difference however, the stationary solution itself also changes when increased costs are associated with E^2, unlike the case of $(dE/dt)^2$ penalties. Another approach would be to focus on viable control instead of only optimal control (Bene et al. [2001], DeLara and Doyen [2008], Krawczyk et al. [2013]) and to seek management strategies that satisfied constraints on fisheries management reflecting the goals of different stakeholders. Indeed, some viable control studies also emphasize political economy challenges involved in implementing fast transitions to desired states (e.g., Martinet et al. [2010]). If taking a viable control approach, managers potentially could be presented with a range of possible management strategies that avoided making fast policy changes provided other desired outcomes of managing a fishery (e.g., ensuring a large NPV of the fishery) could still be met.

Our approach to representing policy adjustment costs was clearly very phenomenological, potentially representing multiple entry-points of policy adjustment costs acting simultaneously. A complementary modeling strategy would be to pick one of those potential entry-points and model how policy adjustments lead to increased costs more mechanistically. However, our understanding of the many possible channels through which policy adjustment costs may manifest themselves remains limited and data to parameterize more complicated models are more limited still. As such, we anticipate use of phenomenological representations of policy adjustment costs, of the type we seek to inform, will continue in natural resource models for the time being.

While our models are highly simplified and phenomenological, we would argue that they provide an entry-point to much broader debates within natural resource management about how well-positioned existing natural resource management institutions are to respond to the dynamics of the ecological systems they are responsible for (Young [2002], Sanchirico and Wilen [2005], Berkes [2007], Armsworth et al. [2015], Bode et al. [2016]). More specifically, our work recognizes that policy makers cannot track ecological dynamics as swiftly as they might like, at least not without incurring significant policy adjustment costs, and examines the implications of this for the resource dynamics and the design of intervention strategies.

Acknowledgments. This work was assisted through participation in the "Pretty Darn Good Control" Working Group at the National Institute for Mathematical and Biological Synthesis, sponsored by the National Science Foundation through NSF Award #DBI-1300426, with additional support from The University of Tennessee, Knoxville. The authors acknowledge helpful discussions and input from working group participants Michael Bode, Carl Boettiger, Iadine Chades, Megan Donahue, Frank Doyle, Marie-Josee Fortin, Alan Hastings, Mandy Karnauskas, Jake LaRiviere, and Claire Paris, and the helpful suggestions from the referees and editor.

ENDNOTE

1. The current value Hamiltonian for our problem is obtained by multiplying the present value Hamiltonian by $e^{\delta t}$ and defining λ_i to be current value multipliers (see Kamien and Schwartz [1991] for further details). A current value Hamiltonian is sometimes favored for applications of this type for developing intuition about the nature of the optimal solution.

REFERENCES

L.G. Anderson [1982], *Optimal Utilization of Fisheries with Increasing Costs of Effort*, Can. J. Fish. Aquat. Sci. **39**, 211–214.

P.R. Armsworth, B.A. Block, J. Eagle, and J.E. Roughgarden [2011], *The Role of Discounting and Dynamics in Determining the Economic Efficiency of Time-Area Closures for Managing Fishery Bycatch*, Theor. Ecol. **4**, 513–526.

P.R. Armsworth, E.R. Larson, S.T. Jackson, D.F. Sax, P. Simonin, B. Blossey, N. Green, M.L. Klein, L. Lester, T.H. Ricketts, M.C. Runge, and M.R. Shaw [2015], *Are Conservation Organizations Configured for Effective Adaptation to Global Change*? Front. Ecol. Environ. **13**, 163–169.

C. Bene, L. Doyen, and D. Gabay [2001], *A Viability Analysis for a Bio-Economic Model*, Ecol Econ. **36**, 385–396.

F. Berkes [2007], *Community-Based Conservation in a Globalized World*, Proc. Natl. Acad. Sci. **104**, 15188–15193.

G. Biais [1995], *An Evaluation of the Policy of Fishery Resources Management by Tacs in European Community Waters from 1983 to 1992*, Aquat. Living Resour. **8**, 241–251.

T. Bjorndal and A. Brasao [2006], *The East Atlantic Bluefin Tuna Fisheries: Stock Collapse or Recovery*? Mar. Resour. Econ. **21**, 193–210.

M. Bode, J.N. Sanchirico, and P.R. Armsworth [2016], *Returns from Matching Management Resolution to Ecological Variation in a Coral Reef Fishery*, Proc. R. Soc B, **283**, 20152828. http://dx.doi.org/10.1098/rspb.2015.2828.

C. Boettiger, M. Mangel, and S. Munch [2015], *Avoiding Tipping Points in Fisheries Management through Gaussian Process Dynamic Programming*, Proc. R. Soc. B. **282**, 20141631.

C. Boettiger, M. Bode, J.N. Sanchirico, J. LaRiviere, A. Hastings, and P.R. Armsworth [2016], *Optimal Management of a Stochastically Varying Population When Policy Adjustment Is Costly*, Ecol. Appl. **26**, 808–817.

G. Brown [1974], *An Optimal Program for Managing Common Property Resources with Congestion Externalities*, J. Pol. Econ. **82**, 163–173.

C. Buskens and H. Maurer [2000], *SQP-Methods for Solving Optimal Control Problems with Control and State Constraints: Adjoint Variables, Sensitivity Analysis and Real-Time Control*, J. Computat. Appl. Math. **120**, 85–108.

R.T. Carson, C.W.J. Granger, J.B.C. Jackson, and W. Schlenker [2009], *Fisheries Management under Cyclical Population Dynamics*, Environ. Resour. Econ. **42**, 379–410.

C. Castilho and P.D.N. Srinivasu [2007], *Bio-Economics of a Renewable Resource in a Seasonally Varying Environment*, Math. Biosci. **205**, 1–18.

C. Clark [1990], *Mathematical Bioeconomics*, John Wiley and Sons Inc., NY.

C.W. Clark and G.R. Munro [1975], *The Economics of Fishing and Modern Capital Theory: A Simplified Approach*, J. Environ. Econ. Manage. **2**(2), 92–106.

C.W. Clark, A.T. Charles, J.R. Beddington, and M. Mangel [1985], *Optimal Capacity Decisions in a Developing Fishery*, Mar. Resour. Econ. **2**(1), 25–53.

J.M. Conrad and M.D. Smith [2012], *Nonspatial and Spatial Models in Bioeconomics*, Nat. Resour. Mode. **25**(1), 52–92.

M. DeLara and L. Doyen [2008], *Sustainable Management of Natural Resources: Mathematical Models and Methods, Environmental Science and Engineering*, Springer-Verlag, Berlin, Heidelberg. [ISBN 9783540790730.]

G. Feichtinger, A. Novak, and F. Wirl [1994], *Limit Cycles in Intertemporal Adjustment Models*, J. Econ. Dyn. Contr. **18**, 353–380.

R. Friedland [2010], *Using Constrained Optimal Control to Model a Socially Sustainable Fishery*, Opt. Contr. Appl. Meth. **31**, 567–580.

F.B. Hanson and D. Ryan [1998], *Optimal Harvesting with Both Population and Price Dynamics*, Math. Biosci. **148**, 129–146.

R.F. Hartl and G. Feichtinger [1987], *A New Sufficient Condition for Most Rapid Approach Paths*, J. Optim. Theory Appl. **54**, 403–411.

S. Jørgensen and P.M. Kort [1993], *Optimal Dynamic Investment Policies under Concave-Convex Adjustment Costs*, J. Econ. Dyn. Contr. **17**(1), 153–180.

M.I. Kamien and N.L. Schwartz [1991], *Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management*, Elsevier, Amsterdam, New York. [ISBN 0.444.01609.0.]

U. Kapaun and M.F. Quaas [2013], *Does the Optimal Size of a Fish Stock Increase with Environmental Uncertainties*? Environ. Resour. Econ. **54**, 293–310.

J.B. Kellner, J.N. Sanchirico, A. Hastings, and P.J. Mumby [2011], *Optimizing for Multiple Species and Multiple Values: Tradeoffs Inherent in Ecosystem-Based Fisheries Management*, Conservat. Lett. **4**, 21–30.

A. Kirsch [1996], *An Introduction to the Mathematical Theory of Inverse Problems*, Springer-Verlag New York, Inc., NY. [ISBN 0-387-94530-X.]

J. B. Krawczyk, A. Pharo, O.S. Serea, and S. Sinclair [2013], *Computation of Viability Kernels: A Case Study of By-Catch Fisheries*, Computat. Manage. Sci. **10**, 365–396.

S. Lenhart and J. Workman [2007], *Optimal Control Applied to Biological Models*, Chapman & Hall/CRC Mathematical and Computational Biology. CRC Mathematical & Computational Biology Series, Boca Raton, Chapman and Hall.

T.R. Lewis [1981], *Evolution of a Renewable Resource under Uncertainty*, Can. J. Econ. **14**, 422–439.

M. Liski, P.M. Kort, and A. Novak [2001], *Increasing Returns and Cycles in Fishing*, Resour. Energy Econ. **23**, 241–258.

D. Ludwig [1980], *Harvesting Strategies for a Randomly Fluctuating Population*, J. du Conseil Int. pour l'Exploration de la Mer. **39**, 168–174.

V. Martinet, O. Thebaud, and A. Rapaport [2010], *Hare or Tortoise? Trade-Offs in Recovering Sustainable Bioeconomic Systems*, Environ. Model. Assess. **15**, 503–517.

B. McGough, A.J. Plantinga, and C. Costello [2009], *Optimally Managing a Stochastic Renewable Resource under General Economic Conditions*, B.E. J. Econ. Anal. Pol. **9**, 56.

National Resource Council [2014], *Evaluating the Effectiveness of Fish Stock Rebuilding Plans in the United States*, The National Academies Press, Washington, DC. [ISBN 978-0-309-29230-6. URL http://www.nap.edu/catalog/18488/evaluating-the-effectiveness-of-fish-stock-rebuilding-plans-in-the-united-states.]

M.G. Neubert [2003], *Marine Reserves and Optimal Harvesting*, Ecol. Lett. **6**(9), 843–849.

M.G. Neubert and G.E. Herrera [2008], *Triple Benefits from Spatial Resource Management*, Theor. Ecol. **1**, 5–12.

L. Nostbakken and J.M. Conrad [2007], Uncertainty in bioeconomic modelling, in (A. Weintraub, ed.), *Handbook of Operations Research in Natural Resources*, *Volume 99 of International Series in Operations Research and Management Science*, Springer, New York, pp. 217–235.

A. Parma [1990], *Optimal Harvesting of Fish Populations with Non-Stationary Stock-Recruitment Relationships*, Nat. Resour. Model. **4**, 39–76.

K. Patterson and M. Resimont [2007], *Change and Stability in Landings: The Responses of Fisheries to Scientific Advice and Tacs*, ICES J. Mar. Sci. **64**, 714–717.

J. Pope [1973], An Investigation into the Effects of Variable Rates of the Exploitation of Fishery Resources, in (M.S. Bartlett and R.W. Hiorns, eds.), *The Mathematical Theory of the Dynamics of Biological Populations*, 347, Academic Press, New York, pp. 23–34.

S. Ralston, A.E. Punt, O.S. Hanel, J.D. DeVore, and R.J. Conser [2011], *A Meta-Analytic Approach to Quantifying Scientific Uncertainty in Stock Assessments*, Fish. Bull. **109**, 217–231.

W.J. Reed [1979], *Optimal Escapement Levels in Stochastic and Deterministic Harvesting Models*, J. Environ. Econ. Manage. **6**(4), 350–363.

J. Roughgarden and P. Armsworth [2001], Ecological Economic Theory for Managing Ecosystem Services, in (M. Press, N. Huntly, and S. Levin, eds.), *Ecology: Achievement and Challenge*, Blackwell Science, Oxford, UK, pp. 337–356.

J.N. Sanchirico and J.E. Wilen [2005], *Optimal Spatial Management of Renewable Resources: Matching Policy Scope to Ecosystem Scale*, J. Environ. Econ. Manage. **50**, 23–46.

J.N. Sanchirico, J.E. Wilen, and C. Coleman [2010], *Optimal Rebuilding of a Metapopulation*, Am. J. Agric. Econ. **4**, 1087–1102.

M. Spence and D. Starrett [1975], *Most Rapid Approach Paths in Accumulation Problems*, Int. Econ. Rev. **16**, 388–403.

O. Tahvonen [2009], *Economics of Harvesting Age-Structured Fish Populations*, J. Environ. Econ. Manage. **58**, 281–299.

F. Wirl [1992], *Cyclical Strategies in Two-Dimensional Optimal Control Models: Necessary Conditions and Existence*, Ann. Operat. Res. **37**, 345–356. [URL http://link.springer.com/article/10.1007/BF02071064.]

O. Young [2002], *The Institutional Dimensions of Environmental Change: Fit, Interplay and Scale*, MIT Press, Cambridge, Mass. [ISBN 9780262240437.]

Supporting Information. Additional Supporting Information may be found in the online version of this article at the publisher's web site:

Figure S1: This plot shows the evolution of the peak-to-valley width ratio for the optimal stock function N^* in the case of time-varying $r(t)$. The peak width is defined as the amount of time the stock spends above the mid-point between the max and the min (below, in the case of valley width.) Note that as c_2 increases, the peak-to-valley ratio approaches one, which is the value it would have for a symmetric sinusoid.

Figure S2: The evolution of the phase lag between valleys of $r(t)$ and valleys of the (a) singular stock and (b) singular effort level in the presence of policy adjustment costs associated with the rate of change of effort. Solutions were calculated for four different values of the regularization parameter (c_2) and a dense set of values for c_1. Note that for each value of the regularization parameter, the general trend is for the phase lag to *decrease* as c_1 increases. (Compare this to Figure 4 c, which shows that for fixed c_1, the phase lag *increases* as c_2 increases.) The regularized phase lags seem to converge toward the unregularized phase lags uniformly in c_1 as $c_2 \to 0$.